



US009460293B1

(12) **United States Patent**
Kashyap et al.

(10) **Patent No.:** **US 9,460,293 B1**
(45) **Date of Patent:** **Oct. 4, 2016**

(54) **CLIPBOARD HARDENING**

(56) **References Cited**

(71) Applicant: **Bromium, Inc.**, Cupertino, CA (US)

U.S. PATENT DOCUMENTS

(72) Inventors: **Rahul Kashyap**, Foster City, CA (US);
Rafal Wojtczuk, Warsaw (PL); **Ian Pratt**, Cambridge (GB)

8,171,390 B1 * 5/2012 Channakeshava G06F 17/243
715/205
2011/0131179 A1 * 6/2011 Ohata G06F 9/468
707/609
2012/0096368 A1 * 4/2012 McDowell G06F 9/543
715/748
2012/0240224 A1 * 9/2012 Payne H04L 63/1416
726/21

(73) Assignee: **Bromium, Inc.**, Cupertino, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 121 days.

* cited by examiner

Primary Examiner — Joseph P Hirl

Assistant Examiner — Chi Nguy

(74) *Attorney, Agent, or Firm* — Brokaw Patent Law PC; Christopher J. Brokaw

(21) Appl. No.: **13/923,212**

(22) Filed: **Jun. 20, 2013**

Related U.S. Application Data

(60) Provisional application No. 61/662,288, filed on Jun. 20, 2012.

(51) **Int. Cl.**
G06F 21/00 (2013.01)
G06F 21/60 (2013.01)

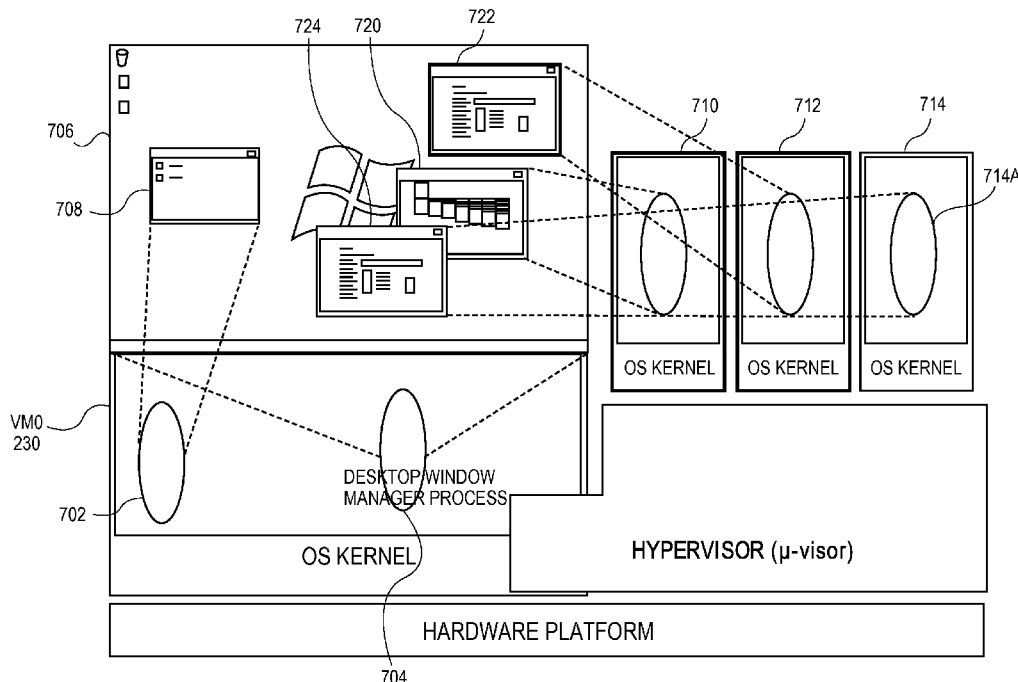
(52) **U.S. Cl.**
CPC **G06F 21/60** (2013.01)

(58) **Field of Classification Search**
USPC 726/1, 26, 30, 31
See application file for complete search history.

(57) **ABSTRACT**

A software module executes on a first operating system running. The software module determines that the first operating system has caused data to be written to a first clipboard maintained by the first operating system. The software module consults policy data to determine whether the data should additionally be written to a second clipboard maintained by a second operating system. The policy data may only allow the data to be written to the second clipboard if the data was written to the first clipboard at the initiation of or approved by a user. If the software module determines that the policy data allows the data to be written to the second clipboard, then the software module, without human instruction, causes the data to be written to the second clipboard.

21 Claims, 9 Drawing Sheets



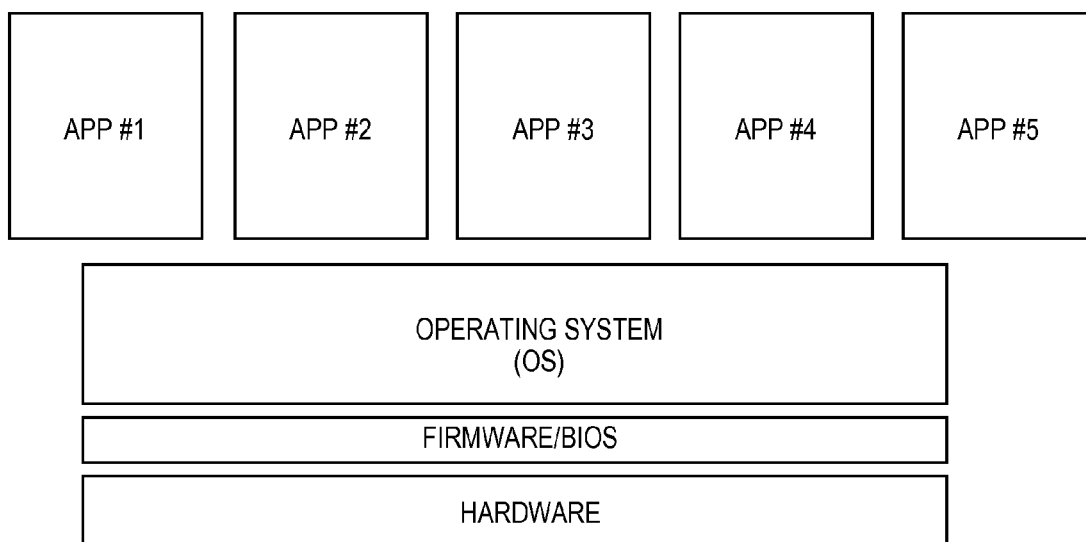


FIG. 1
(PRIOR ART)

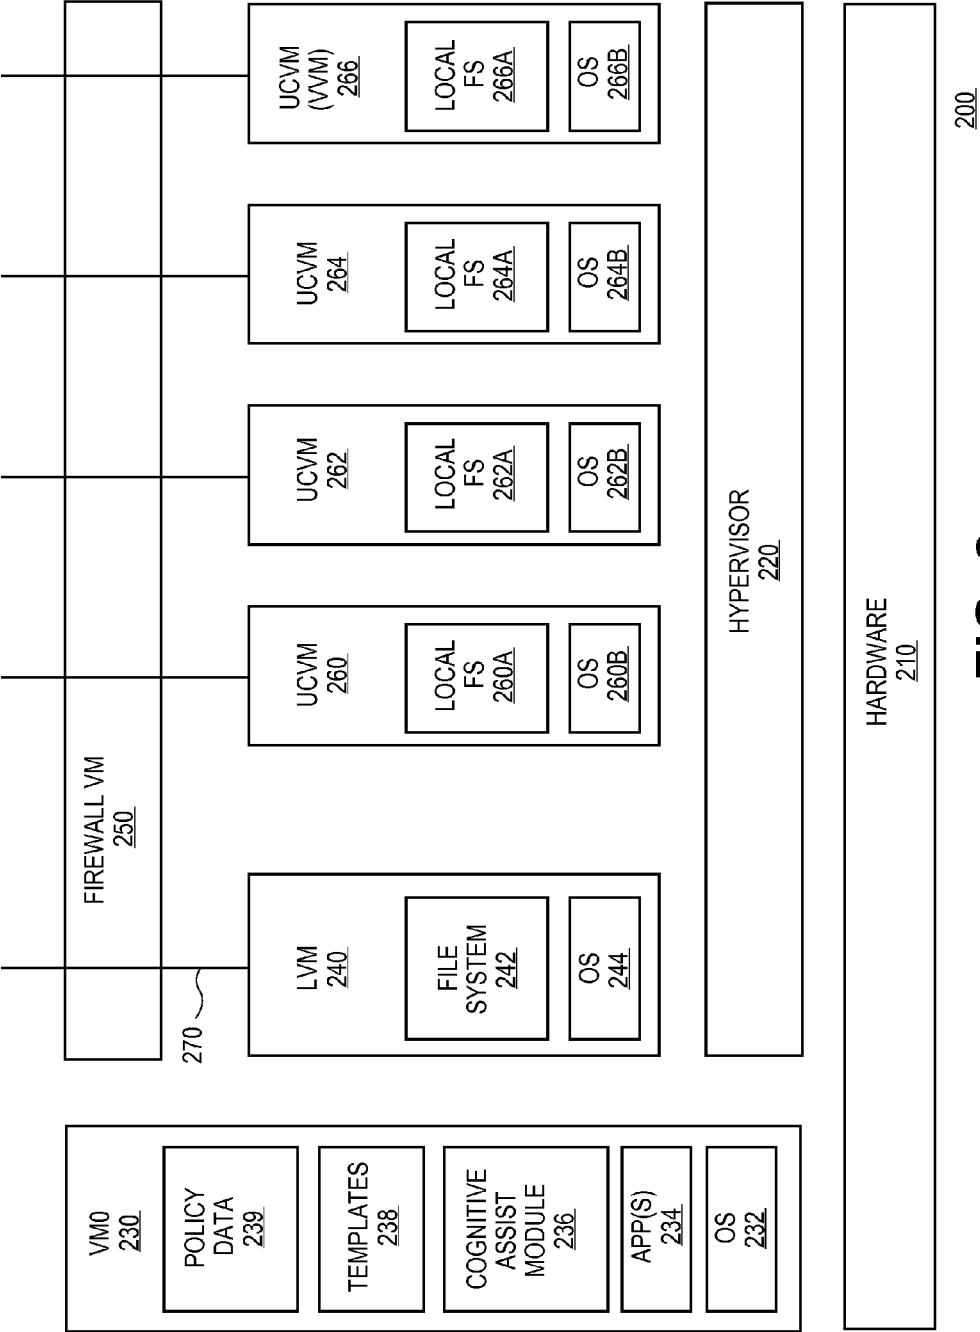


FIG. 2

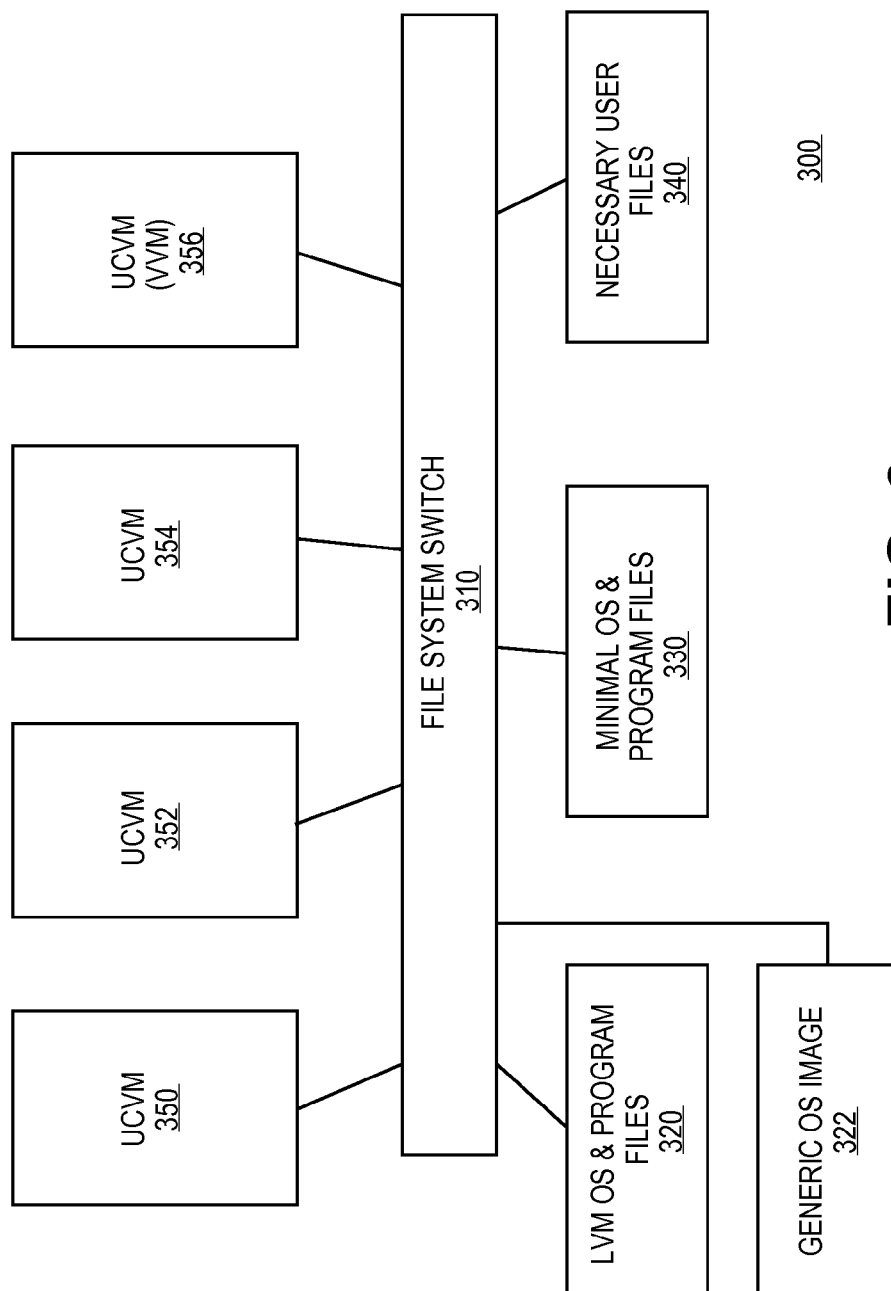
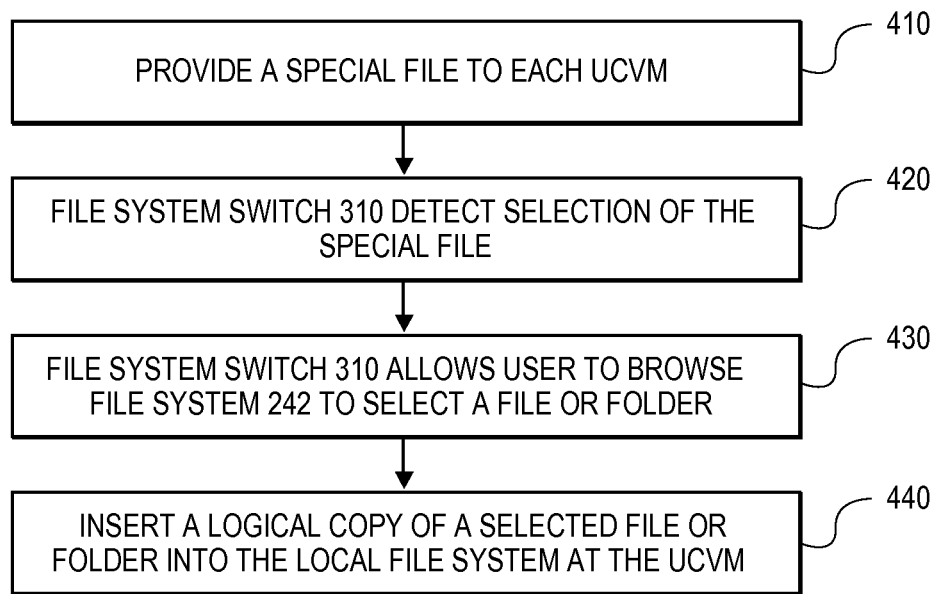


FIG. 3

**FIG. 4**

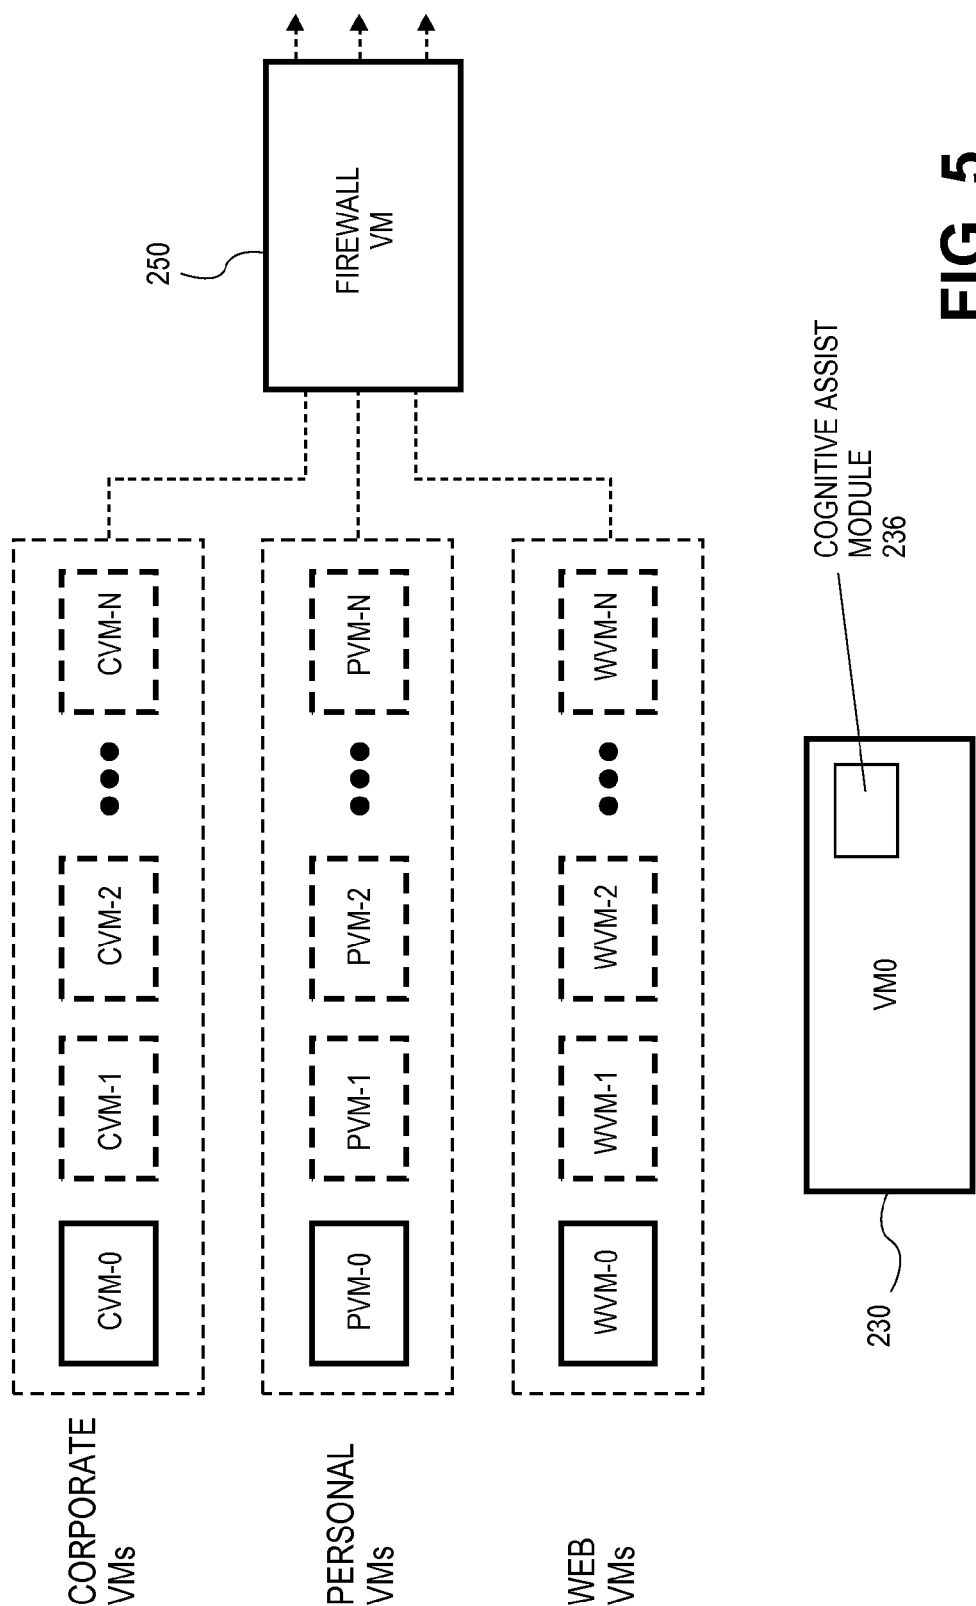
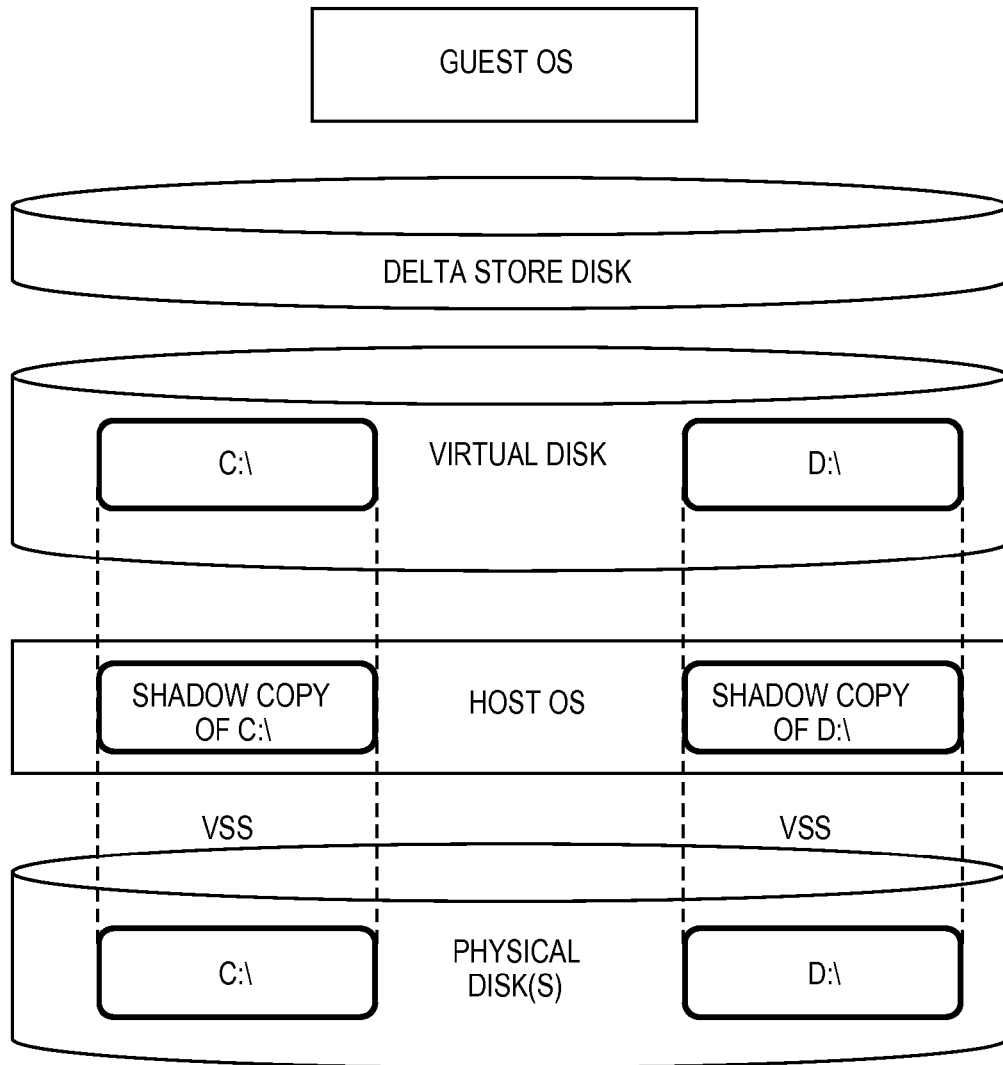


FIG. 5

**FIG. 6**

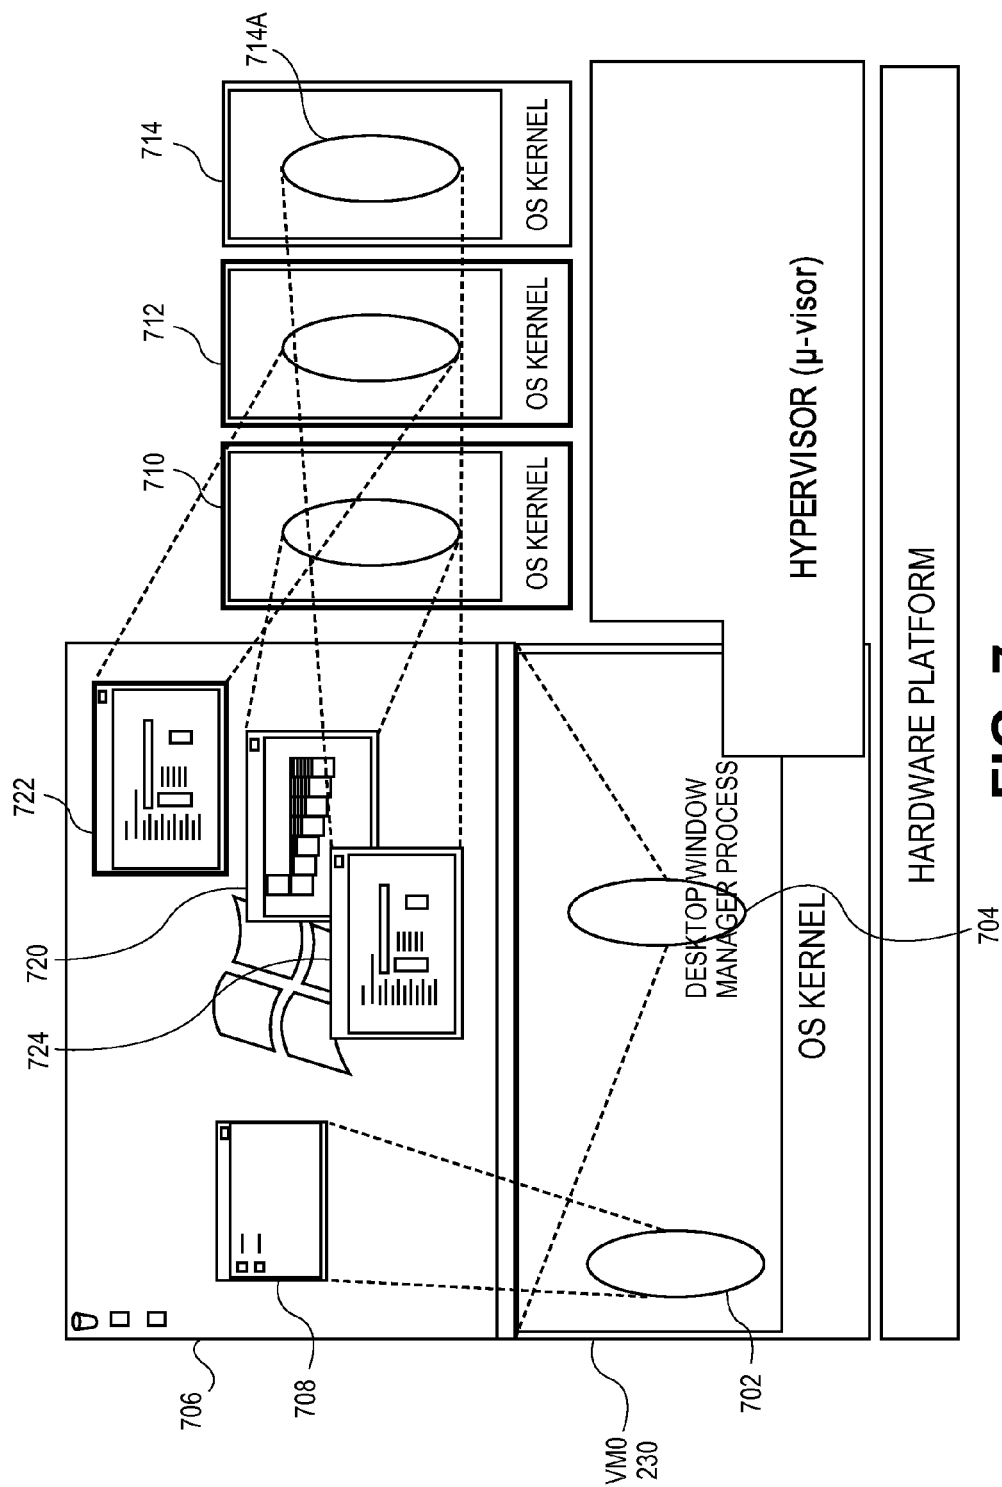


FIG. 7

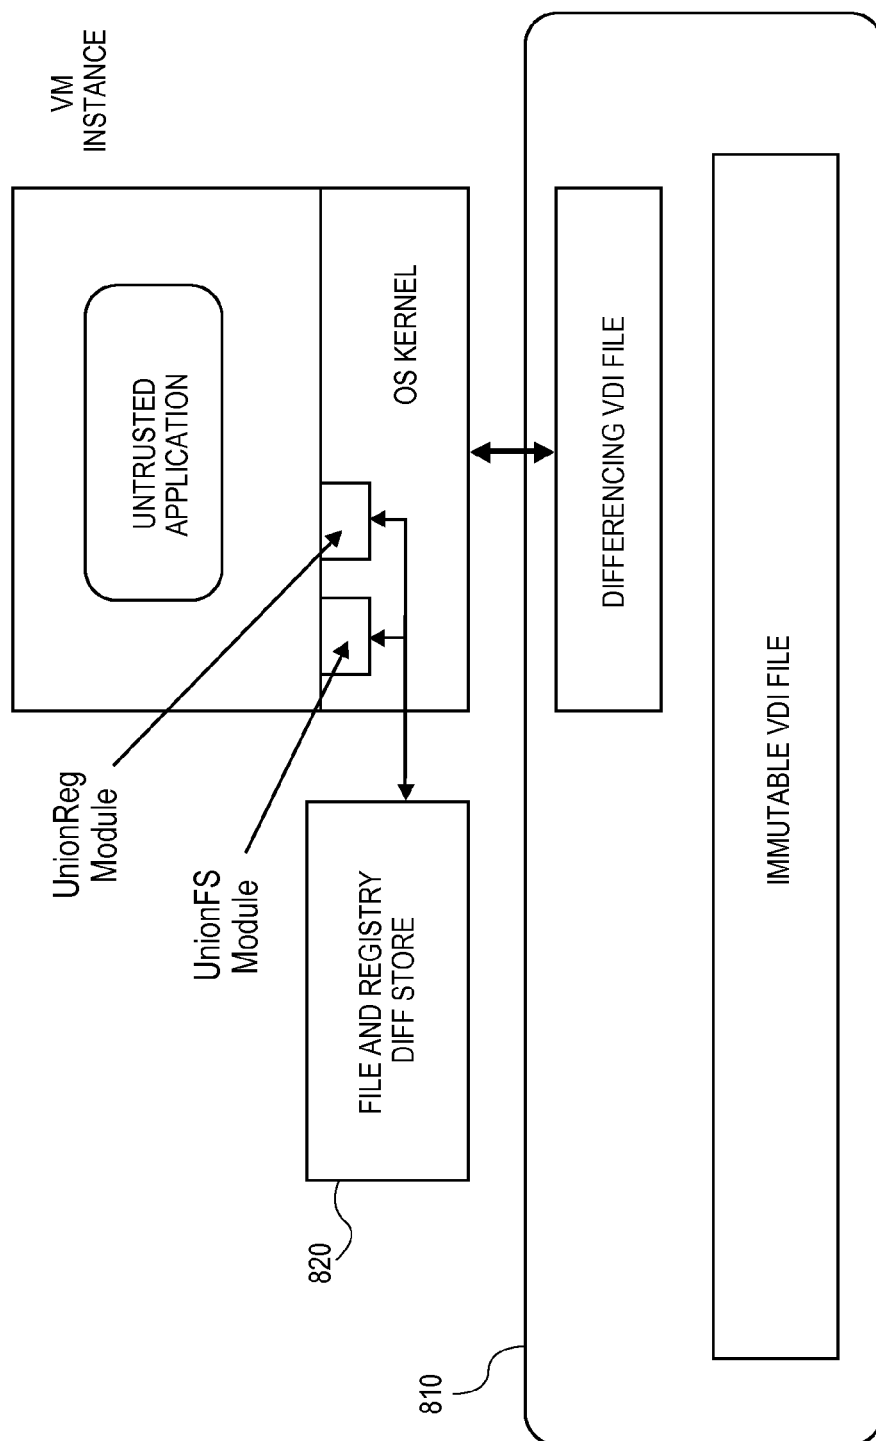
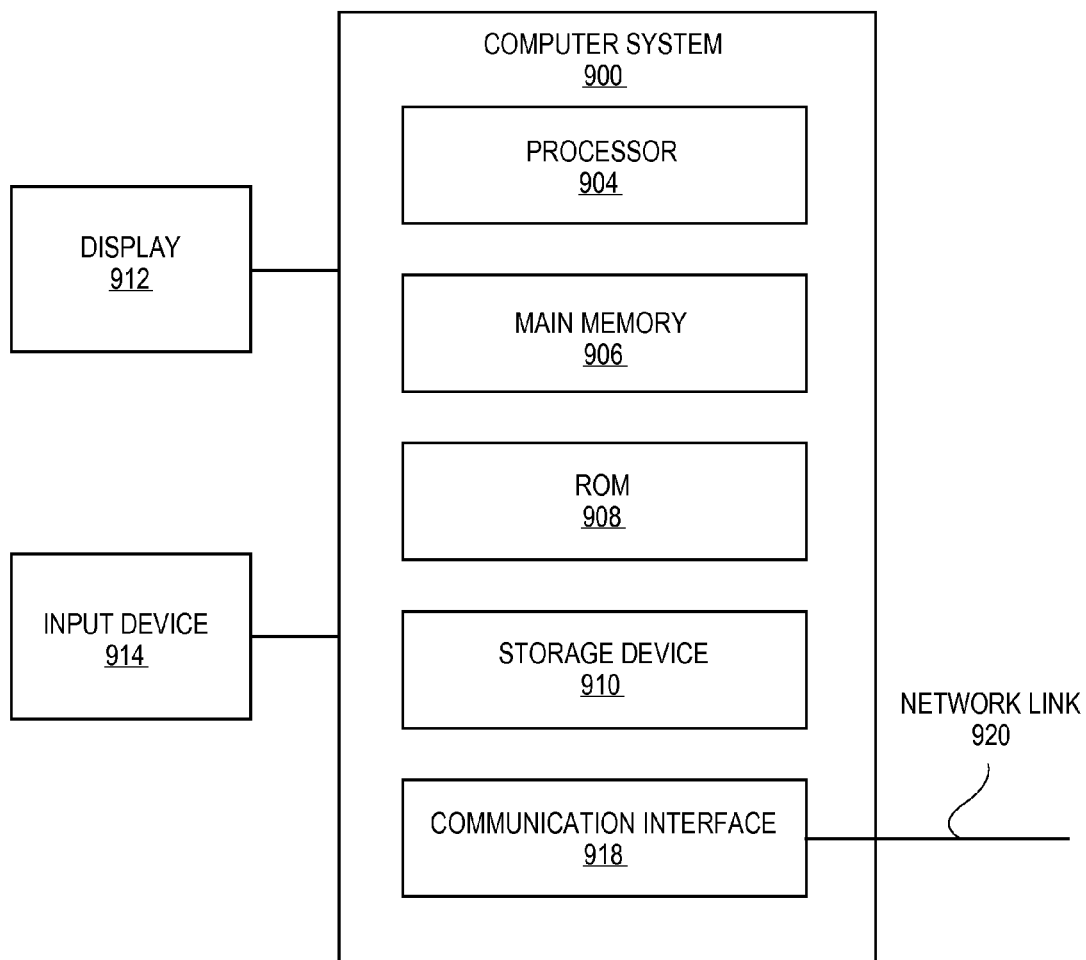


FIG. 8

**FIG. 9**

CLIPBOARD HARDENING**CLAIM OF PRIORITY AND RELATED
APPLICATION DATA**

This application is a U.S. non-provisional patent application, and claims priority to, U.S. provisional patent application No. 61/662,288, filed Jun. 20, 2012, entitled "Clipboard Hardening Across Virtual Machines," invented by Rahul Kashyap et al., the disclosure of which is hereby incorporated by reference in its entirety for all purposes as if fully set forth herein.

This application is related to U.S. patent application Ser. No. 13/115,354, filed on May 25, 2011, entitled "Approaches for Securing an Internet Endpoint Using Fine-Grained operating System Virtualization," invented by Gaurav Banga et al., the disclosure of which is hereby incorporated by reference in its entirety for all purposes as if fully set forth herein.

This application is related to U.S. patent application Ser. No. 13/223,091, filed on Aug. 31, 2011, entitled "Approaches for Automated Management of Virtual Machines for Running Untrusted Code Safely," invented by Gaurav Banga et al., the disclosure of which is hereby incorporated by reference in its entirety for all purposes as if fully set forth herein.

This application is related to U.S. patent application Ser. No. 13/302,123, filed on Nov. 22, 2011, entitled "Approaches for Efficient Physical to Virtual Disk Conversion," invented by Gaurav Banga et al., the disclosure of which is hereby incorporated by reference in its entirety for all purposes as if fully set forth herein.

This application is related to U.S. patent application Ser. No. 13/419,345, filed on Mar. 13, 2012, entitled "Seamless Management of Untrusted Data Using Virtual Machines," invented by Gaurav Banga et al., the disclosure of which is hereby incorporated by reference in its entirety for all purposes as if fully set forth herein.

This application is related to U.S. patent application Ser. No. 13/358,434, filed on Jan. 25, 2012, entitled "Approaches for Protecting Sensitive Data Within a Guest Operating System," invented by Gianni Tedesco et al., the disclosure of which is hereby incorporated by reference in its entirety for all purposes as if fully set forth herein.

This application is related to U.S. patent application Ser. No. 13/468,742, filed on May 10, 2012, entitled "Reducing Memory Requirements for Guest Operating Systems," invented by Krzysztof Uchonski et al., the disclosure of which is hereby incorporated by reference in its entirety for all purposes as if fully set forth herein.

This application is related to U.S. patent application Ser. No. 13/526,409, filed on Jun. 18, 2012, entitled "Synchronizing Resources of a Virtualized Browser," invented by Kiran Bondalapati et al., the disclosure of which is hereby incorporated by reference in its entirety for all purposes as if fully set forth herein.

This application is related to U.S. patent application Ser. No. 13/526,403, filed on Jun. 18, 2012, entitled "Synchronizing History Data Across a Virtualized Web Browser," invented by Kiran Bondalapati et al., the disclosure of which is hereby incorporated by reference in its entirety for all purposes as if fully set forth herein.

This application is related to U.S. patent application Ser. No. 13/526,390, filed on Jun. 18, 2012, entitled "Synchronizing Cookie Data Using a Virtualized Browser," invented by Prakash Buddhiraj et al., the disclosure of which is

hereby incorporated by reference in its entirety for all purposes as if fully set forth herein.

This application is related to U.S. patent application Ser. No. 13/526,373, filed on Jun. 18, 2012, entitled "Composing the Display of a Virtualized Web Browser," invented by Prakash Buddhiraj et al., the disclosure of which is hereby incorporated by reference in its entirety for all purposes as if fully set forth herein.

This application is related to U.S. patent application Ser. No. 13/526,354, filed on Jun. 18, 2012, entitled "Transferring Files Using a Virtualized Application," invented by Deepak Khajuria et al., the disclosure of which is hereby incorporated by reference in its entirety for all purposes as if fully set forth herein.

This application is related to U.S. patent application Ser. No. 13/526,755, filed on Jun. 19, 2012, entitled "Ensuring the Privacy and Integrity of a Hypervisor," invented by Ian Pratt, the disclosure of which is hereby incorporated by reference in its entirety for all purposes as if fully set forth herein.

This application is related to U.S. patent application Ser. No. 13/526,423, filed on Jun. 18, 2012, entitled "Securing Resources of a Virtual Machine," invented by Deepak Khajuria et al., the disclosure of which is hereby incorporated by reference in its entirety for all purposes as if fully set forth herein.

This application is related to U.S. patent application Ser. No. 13/419,356, filed on Mar. 13, 2012, entitled "Securing File Trust with File Format Conversions," invented by Gaurav Banga et al., the disclosure of which is hereby incorporated by reference in its entirety for all purposes as if fully set forth herein.

This application is related to U.S. patent application Ser. No. 13/468,781, filed on May 10, 2012, entitled "Preparing a Virtual Machine for Template Creation," invented by Krzysztof Uchonski et al., the disclosure of which is hereby incorporated by reference in its entirety for all purposes as if fully set forth herein.

This application is related to U.S. patent application Ser. No. 13/468,805, filed on May 10, 2012, 2012, entitled "Managing the Eviction Process," invented by Krzysztof Uchonski et al., the disclosure of which is hereby incorporated by reference in its entirety for all purposes as if fully set forth herein.

This application is related to U.S. patent application Ser. No. 13/468,843, filed on May 10, 2012, entitled "Mitigating Eviction by Maintaining Mapping Tables," invented by Krzysztof Uchonski et al., the disclosure of which is hereby incorporated by reference in its entirety for all purposes as if fully set forth herein.

This application is related to U.S. patent application Ser. No. 13/523,859, filed on Jun. 14, 2012, entitled "Securing an Endpoint by Proxying Document Object Models and Windows," invented by Gaurav Banga et al., the disclosure of which is hereby incorporated by reference in its entirety for all purposes as if fully set forth herein.

This application is related to U.S. patent application Ser. No. 13/526,769, filed on Jun. 19, 2012, entitled "Securing a Controller of a Device from a Host Operating System," invented by Ian Pratt, the disclosure of which is hereby incorporated by reference in its entirety for all purposes as if fully set forth herein.

This application is related to U.S. patent application Ser. No. 13/526,810, filed on Jun. 19, 2012, entitled "Safe Printing," invented by Deepak Khajuria et al., the disclosure of which is hereby incorporated by reference in its entirety for all purposes as if fully set forth herein.

This application is related to U.S. patent application Ser. No. 13/527,161, filed on Jun. 19, 2012, entitled "Network Partitioning," invented by Xin Li et al., the disclosure of which is hereby incorporated by reference in its entirety for all purposes as if fully set forth herein.

FIELD OF THE INVENTION

Embodiments of the invention relate to ensuring virtualized environments share resources responsibly.

BACKGROUND

Ensuring the security of Internet users and Internet connected devices is one of the grand challenges facing us today. The current state of affairs is very problematic, as our cyber-security infrastructure is easily and routinely subverted by cyber criminals, resulting in great economic loss. Every year brings deeper and more complex dependence by society on our cyber-infrastructure, and yet at the same time the cyber-security problem only worsens as the capabilities of the cyber-criminal mature. In effect, we are building mission-critical dependence into virtually every aspect of human activities on a cyber-infrastructure that is very insecure at its core.

The current state of our cyber-security infrastructure is due, at least in part, to two fundamental limitations. The first limitation is a fundamental mismatch between the design assumptions made by computer security programmers with how the vast majority of users interact with the cyber-infrastructure (the "Security Model Complexity" problem). The second limitation is a lack of appropriate isolation of code and data from trusted and untrusted sources in modern computer systems (the "Lack of Isolation" problem). These two limitations of current systems are somewhat orthogonal, but are both very important for securing a computer system. The "Lack of Isolation" problem, in particular, is very important because modern computer devices, such as a PC or cell phone, are typically general purpose devices that execute a wide variety of software from different sources.

The general purpose capability of modern computing systems is constructed using a layered stack of hardware and software. An example of the layered arrangement of hardware and software that is present in modern computer systems is shown in FIG. 1. At the lowest layer, there is hardware with a small number of basic general purpose programming capabilities. Upon this hardware layer sits the firmware/BIOS which is responsible for, among other things, initializing hardware resources and loading the operating system. The operating system (OS) provides a file system and functionality which may be used by a variety of different applications. On top of the OS layer run the various applications which provide user-visible rich functionality to the computer. The functionality provided by the application layer is typically the primary concern of the computer user.

One advantage and consequence of the layered nature of modern computer systems is that the various layers may come from different vendors, as long as the layers conform to the specifications governing the layer boundary (which may be based on open or proprietary industry standards). To illustrate an example, in a typical PC today the hardware may be constructed around processor and chipset technology provided by Intel or AMD. The firmware/BIOS may be provided by companies like Insyde, AMI or Phoenix Technologies and may be written to conform to several industry specifications such as UEFI and PI. The operating system (OS) may originate from a company like Microsoft or Apple

or may be a flavor of the Linux open source OS. Finally, the applications themselves are usually written to the specification of one of the operating systems and may be provided by one of a large multitude of application vendors or obtained from a variety of different sources.

Note that some of the applications may themselves have a layered architecture. A web browser, for example, typically includes a browser core and may also download web applications in the form of HTML, JavaScript and Flash programs from various Internet web sites. The web browser may run these downloaded web applications locally on top of the browser core. A typical web page contains HTML with embedded JavaScript that can change the HTML being rendered by the web browser dynamically based on user actions without having to re-download the web page from the web server. The HTML may also demarcate part of the web page to be rendered by a plug-in, which is typically a separate program that is installed on the computer. Plug-ins are often downloaded from different sources over the World Wide Web. Thus, a modern computer runs code that comes from a variety of different sources. In particular, application programs may originate from literally millions of different sources once we consider the collection of traditional local applications as well as web applications that are downloaded from web sites.

The integrity of a computer system when it runs application code from different sources (or even the same program being run by different users of a shared computer) has traditionally been one of the responsibilities of the OS. The OS uses various hardware and software constructs like virtual memory, processes, and file permissions to prevent programs belonging to one user from affecting programs (and data) belonging to another user. This responsibility of the OS to "isolate" programs and data from one another often tends to compete with another responsibility of the OS, which is to allow for cooperation between programs especially between user application programs and system level services such as shared library modules, database services, and other higher-level common OS functionality. These two OS functions, to share and to isolate, require the OS designer to make certain tradeoffs on how much to share and how much to isolate.

As a result of these design tradeoffs, the resulting implementation of modern operating systems has grown to a complexity such that ensuring the OS has no security issues is impractical. In mature operating systems, the security implementation is typically robust enough to work well for normal programs under normal usage with no adverse impact on the operation of the computer. However, most OS implementations are very large and complex bodies of computer code that may not possess a sufficiently robust security system when interacting with programs that are especially designed to take advantage of less-tested or unvalidated corner cases in the operation of the security subsystem.

These "security vulnerabilities" are not important for well-behaved programs during typical operation, but are used extensively by cyber criminals to subvert the computer's security subsystems. Once the system's security is subverted, it is generally possible for cyber criminals to run any software under their control on the subverted computer system.

The Lack of Isolation problem stems from the fact that a large amount of code executed by computers today comes from sources outside the computer, some of which have explicit intentions of committing criminal activities. This includes any program downloaded from the Internet or any

web site visited by the computer. All downloaded programs (good and bad) have the same OS and library services available to them to use during their operation. Consequently, any program (even malware), can exploit any security vulnerability in the complex OS or web browser environment and subvert the security subsystem that isolates applications from one another. For example, when a user visits a web site, he or she is really running web application code developed by the publisher of the web site. If this web site is malicious, then malware may be executed on the computer. Malware may be designed to exploit a security vulnerability in the web browser to take control of the computer system during subsequent web site visits, e.g., if you visit your bank's web site, your key strokes may be captured and your login/password information for the bank may be transmitted to the malware publisher. Malicious programs may be designed for a variety of purposes, e.g., a malicious program may simply be designed to interfere with the normal operation of a computer rather than extracting useful data from the computer.

While some computer security professionals may understand the existence of the Lack of Isolation problem, this problem is hard to solve in any practical way because preventing applications from working or communicating with each other tends to conflict with achieving the goal of increasing seamless communication between different local and web applications. There has been some work towards the isolation of web code from different sources being run by a web browser. Modern browsers have attempted to create a level of sandboxing around downloaded web application code in order to isolate downloaded code from the rest of the computer and from each other. However, these models are fairly primitive in their ability to deal with the full gamut of security issues that arise during the course of a typical user's web experience. For example, certain versions of Google's Chrome web browser's sandboxing does not address safety issues arising from downloaded browser plug-ins and various types of native executables; thus, every computer system running certain versions of Chrome is vulnerable to a zero day exploit attack against Adobe Flash or Microsoft Word as much as if the system was running a less secure or older browser with the same Adobe Flash Plug-in or Microsoft Word plug-in.

Web browsers have been burdened with the need to ensure full compatibility to older and non-standard web pages in their efforts to provide superior safety and privacy. For example, web browser programmers have had to make some relaxations in order to correctly render popular web sites that rely on the sharing of information between web sites.

Last but not least, most web browsers vendors suffer from a huge conflict of interest because their business relies upon monetizing the web browsing habits of their users within their own business processes and with their industry partners. This monetization relies on data about users' browsing habits which is contained in the web cookies that are set and later provided to web servers during the course of web sessions. Companies such as Google and Microsoft have a great interest in learning as much as possible about a person's browsing habits and typically arrange the default privacy settings of web browsers to be advantageous to them (but less than optimal from a security and privacy standpoint). This choice of default privacy and core functionality settings causes web browsers to transfer large amounts of sensitive information from end users' machines to Internet related businesses, such as Google, Microsoft, Apple, etc., thereby allowing such businesses to better monetize their customer base by offering appropriate products and services

and serving targeted ads. These same settings, however, can be leveraged by malicious parties to exploit security vulnerabilities. While all web browsers provide some level of control to the sophisticated user to tune his or her web browser functionality and/or privacy/safety settings to browse more securely, the vast majority of users never change these default settings.

Some security researchers have also proposed the use of "client virtualization" (also called "Virtualization using a Hypervisor" in the desktop) to solve the Lack of Isolation Problem. In one form of client virtualization, the user runs multiple independent operating systems on their laptop or desktop on multiple virtual machines (VMs) within the client system which have been created using a hypervisor, such as from VMWare of Palo Alto, Calif. or Virtual PC, available from Microsoft Corporation of Redmond, Wash. When client virtualization is used to achieve improved security, different VMs are used to run applications from different sources or of different types. For example, an OS in one VM may be dedicated for accessing the corporate network that the user may be part of and running corporate applications (local and web). Another OS in a second VM might be used by the user to run his or her personal programs and store personal documents. Finally, a different OS in a third VM may be used for general web browsing on the wider Internet and running native executables that may have been downloaded from the Internet. An example of such a solution is XenClient, which is made by Citrix Systems of Ft Lauderdale, Fla.

The use of classical client virtualization, as discussed above, to solve the general code isolation problem in the context of Internet endpoints suffers from several drawbacks. A first drawback is that there is too much management overhead for the end-user. The end-user has the onus of making the decision as to what VM to use for each activity. Any mistake, intentional or accidental, may subvert the integrity of the system. While many safeguards can be added as a layer on top of the core virtualization technology to help prevent the user from making mistakes, this has not yet been demonstrated to work in a practical and robust fashion.

An additional drawback is that client virtualization, as described above, suffers from the problem that any VM that is used for general web browsing is just as vulnerable to a security problem as any monolithic system running a single VM while accessing web sites on the general Internet. Therefore, it is quite likely that the VM dedicated to web browsing described in the arrangement above will be subverted by malware eventually. Any subsequent activities in that VM, then, will be compromised.

BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments of the invention are illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

FIG. 1 is an illustration of the layered arrangement of hardware and software present in modern computer systems;

FIG. 2 is an block diagram of the functional components of one embodiment of the invention;

FIG. 3 is block diagram of the functional components involved in exposing a restricted copy of the file system to different UCVMs (and VVMs) according to an embodiment of the invention;

7

FIG. 4 is a flowchart illustrating the steps involved in a UCVm obtaining a copy of a new user file maintained in the file system stored elsewhere according to an embodiment of the invention;

FIG. 5 is an illustration of instantiating a plurality of different virtual machines using different templates according to an embodiment of the invention;

FIG. 6 is an illustration of a virtual disk based on VSS shadow copies according to an embodiment of the invention;

FIG. 7 is an illustration of exemplary desktop of a client according to an embodiment of the invention;

FIG. 8 is an illustration of safely installing an untrusted application according to an embodiment of the invention of the invention; and

FIG. 9 is a block diagram that illustrates a computer system upon which an embodiment of the invention may be implemented.

DETAILED DESCRIPTION OF THE INVENTION

Approaches for securely managing objects copied to the clipboard in a virtualized environment. In the following description, for the purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the embodiments of the invention described herein. It will be apparent, however, that the embodiments of the invention described herein may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form or discussed at a high level in order to avoid unnecessarily obscuring teachings of embodiments of the invention.

Client Overview

Embodiments of the invention prevent malicious code, introduced into a computer system, from compromising the resources of the computer system through the use of dynamic operating system (OS) micro-virtualization. A computer system of an embodiment includes a number of independent virtual machines (VMs) that each executes a full operating system (OS). A block diagram of client 200 according to one embodiment of the invention is shown in FIG. 2. The term “client,” as broadly used herein, represents any type of Internet endpoint or computer system capable of connecting to a network and executing a virtual machine. Non-limiting, illustrative examples of client 200 include a PC, a laptop computer, a tablet computer, a cell phone, a personal digital assistant (PDA), and the like.

In an embodiment, client 200 may correspond to a server. Thus, while use of the term “client” in other contexts might exclude an interpretation that includes a server, as broadly used herein, client 200 may be embodied on a wide variety of machines, one example of such being a server. Thus, as the Applicant may be his or her own lexicographer, as used herein, the term client 200 expressly includes a server. For example, non-limiting, illustrative examples of client 200 include a web server, an application server, a file server, and a cloud server. Indeed, implementing embodiments of the invention upon a server may yield many benefits. The micro-virtualization techniques employed by embodiments provide an efficient mechanism for eliminating the risk of executing untrusted code and/or interpreting untrusted data in accordance with different policies to manage such risk. As such, a device, such as a server, which interacts with (a) numerous sources of untrusted code and/or data and/or (b)

8

two or more corporate entities having different policies towards managing the risk of untrusted code and/or data, may benefit from embodiments of the invention.

Client 200 includes a number of virtual machines (such as 230, 240, 250, and 260, for example) that execute on hardware 210 of client 200. The various VMs within client 200 may be used for separately executing processes associated with different activities. One such VM, namely “VM0” (i.e., VM0 230 of FIG. 2), is secured so that VM0 may serve as the root of trust with a guaranteed integrity. VM0 may contain core operating system 232 and one or more applications 234. In the embodiment shown in FIG. 2, VM0 is not accessible over any network, such as the Internet. As shall be explained below, VM0 provides a secure environment in which operating system 232 and one or more applications 234 may execute without risk of exposure to malicious code.

Other VMs, such as VMs 260, 262, 264, and 266 in FIG. 2, may be created, maintained, and destroyed on-demand using a very efficient micro-virtualizing hypervisor 220. Using efficient micro-virtualization techniques, the latency of starting and stopping activities or applications which run in their own VM in embodiments is very low, thereby providing a practical user experience while employing full OS virtualization.

Embodiments address and overcome many disadvantages, such as the Lack of Isolation Problem, experienced by modern general purpose computer systems that execute code from different sources and of differing trust levels; nevertheless, embodiments maintain compatibility with current typical real-world usage of computer systems by corporate and non-corporate users. This is so because any activity which is not previously deemed trustworthy is performed in a separate VM by certain embodiments, and so all code which may be potentially malicious is executed in its own VM that is destroyed after its immediate use is ended, thereby preventing any malicious code from effecting any lasting change to a computer system according to an embodiment of the invention.

The Trusted Virtual Machine—Vm0

In an embodiment of the invention, a special virtual machine, referred to herein as “VM0,” is created to be a trusted and un-hackable portion of a computer system. FIG. 2 depicts VM0 230 according to an embodiment. To achieve the property of being un-hackable, VM0 230 may be permanently disconnected from any network (i.e., VM0 230 is not connected to any local network or the Internet). Specifically, VM0 230 may not contain any type of networking stack, such as a TCP/IP network stack, and may not have access to any networking hardware that could allow for communication between VM0 230 or any applications 234 executed thereby and the Internet. Thus, the only way to install software onto VM0 230 is to have physical custody of client 200 and manually install the software on VM0 230.

Thus, in certain embodiments, one or more applications 234 executing within VM0 230 do not have any access to a network, must be fully self contained in their functionality, and must rely only on local code and data for all their functionality. All applications that need to access the network will therefore need to run in a separate virtual machine outside of VM0 230, as shall be described in further detail below. It is envisioned that the software (such as one or more applications 234) running in VM0 230 be selected at the time client 200 is manufactured or first configured for use in a controlled environment. Because VM0 230 is never con-

nected to any type of network, such as a TCP/IP network, all common types of network initiated attacks cannot be waged against VM0 230, thereby rendering VM0 230 immune to such attacks and safe as compared to any computer or VM that is connected to the Internet.

In an embodiment where hypervisor 220 is a Type 2 hypervisor, when client 200 is booted, only VM0 230 is started by the BIOS or firmware of client 200. Once VM0 230 is running, VM0 230 can start hypervisor 220 immediately or on demand. In another embodiment, where hypervisor 220 is a type 1 hypervisor, hypervisor 220 is first started by the BIOS when client 200 is booted and VM0 230 is launched by the Hypervisor 220. Hypervisor 220 is a software component that is responsible for creating other VMs which each execute independent instances of the operating system. These additional VMs are instantiated by VM0 230 and/or hypervisor 220 to run any untrusted code or code that needs to access the network. Untrusted code in this context is any code which has not been pre-approved as being trusted by an IT administrator of client 200. The additional VMs are started “silently” and automatically by client 200, e.g., these VMs are started transparently to the user and without the user having to do anything explicit. These additional VMs are also not explicitly visible to the user; instead, all the user sees on the desktop is familiar objects (such as icons, windows, and applications) without any indication that multiple VMs are executing in client 200. Embodiments of the invention follow rules that govern what application activities are assigned to which particular VM. These rules are described below in greater detail.

In another embodiment (not depicted in FIG. 2), VM0 230 may have a networking stack that is firewalled off from the network using well-tested firewall software, thereby allowing VM0 230 to have access to a computer network. Such an embodiment may only allow connections with a specific Internet system so that the software inside VM0 230 may be updated from a designated update server. For example, the firewall software may only allow VM0 230 to connect to one or more servers associated with the IT administrator of client 200 and may prevent VM0 230 from establishing a connection with any other endpoint on any network.

Interaction with an User Interface

All code responsible for generating a user interface (UI) not associated with an application may be maintained in VM0 230. Consequently, all UI interaction activity with the desktop between a user and software executing on client 200 may take place between the user and VM0 230, which maintains a unified desktop for all applications running in all VMs. Interaction between the user and applications running in VMs other than VM0 230 takes place indirectly via VM0 230. For example, when the user enters a password for a web site whose browser and HTML/Javascript code is running in an untrusted VM, the password is first directly provided to VM0 230, which then transfers the information to the untrusted VM. Furthermore, the untrusted VM's display is rendered on to a virtualized display, which is then composed into the VM0 230 desktop (as appropriate) by controlling code running in VM0 230. As code executing in VM0 230 is trusted, the user may trust any user interface controls displayed on a screen since all code responsible for rendering the user interface is trusted.

This approach is quite different from prior systems where often the code that controls the full desktop experience is untrusted. Consequently, if the code responsible for generating the user interface is corrupted by malware, then the

user interface may be used as a tool to deceive the user. For example, malware may cause a user interface control to be displayed that requests the user to submit an authentication credential which will be used for improper purposes by the malware. However, this problem is overcome by embodiments of the invention—since all code responsible for rendering user interface controls executes in VM0 in an embodiment, malware is prevented from hijacking or corrupting UI-rendering code.

To illustrate an embodiment of the invention, consider FIG. 7, which is an illustration of exemplary desktop of client 200 according to an embodiment. As shown in FIG. 7, process 704 is responsible for rendering desktop 706 on a physical display of client 200. Process 714A runs in untrusted VM 714 and does not have complete access to the file system of client 200. When any process inside VM 714 requests access to the file system of client 200, it is intercepted and process 702 is responsible for rendering a window 708 depicting the contents of the file system of client 200. Process 702 has the option of selectively displaying which contents are available to the VM 714 based on policies as set forth by the IT administrator or the user. VM 710 in FIG. 7 that runs the solitaire game is implemented such that the display of VM 710 is a virtualized display, which is then composed into the desktop 706 (as appropriate) by controlling process 704 running in VM0 230. The displays of VMs 712 and 714 are rendered on the desktop 706 in a similar fashion.

The Legacy Virtual Machine—LVM

FIG. 2 depicts a legacy virtual machine (LVM) 240 according to an embodiment of the invention. LVM 240 may contain operating system 244. LVM 240 serves as the primary entity being managed by the IT administrator of client 200. As such, LVM 240 provides an environment that is analogous to the managed enterprise OS of corporate computer system in that an IT department may install and maintain various enterprise applications within operating system 244 of LVM 240. In an embodiment, operating system 244 of LVM 240 may correspond to a Microsoft Windows OS or any other general purpose OS such as Linux or MacOS.

In an embodiment, LVM 240 is responsible for storing the main file system 242 of client 200. File system 242 may contain the user's profile folder containing the user's settings and files.

LVM 240 typically only runs infrastructure OS programs and programs that are used for the purpose of managing client 200 and trusted enterprise applications. Other user programs (especially those that involve external components or consume untrusted data) do not run in LVM 240, but instead, run elsewhere in separate VMs (such as a UCVM as described in more detail below).

In an embodiment, the network access of LVM 240 is restricted to just the corporate network as implemented by firewall VM 250. Firewall VM 250 is a specialized virtual machine that comprises firewall software/applications to restrict network access of VMs running in client 200 to appropriate and/or necessary network access points. Such practice is consistent with the need for only the responsible IT administrator to be capable of connecting to LVM 240 to manage LVM 240 and processes executing therein.

In one embodiment, LVM 240 and VM0 230 may be implemented in a single virtual machine.

Untrusted Code Virtual Machine—UCVM

When a user wishes to run any application that requires access to either a network or untrusted data (untrusted data

11

is any data that originates from outside client **200**), the application is run inside a dedicated VM that is created on-demand by hypervisor **220**. This dedicated VM is called an Untrusted Code Virtual Machine (or UCVm). FIG. 2 depicts several UCVms, namely UCVm **260**, **262**, **264**, and **266**. A UCVm operates under the assumption that, in general, any code that connects to the network and interacts with arbitrary code executing on an external device may at some point be compromised. This assumption also applies to trusted applications that interact with data originating from outside the computer system executing the trusted application, because such data may, at some point, contain embedded malicious code. To address such possibilities, such applications are executed in a UCVm to prevent any malicious code, inadvertently introduced into the UCVm, from having the capacity to affect any change outside of the UCVm.

In an embodiment, a UCVm is created by (a) cloning a copy of LVM **240**, or a stripped-down version of LVM **240**, in memory and (b) providing access to a restricted file system to the newly created UCVm. For example, UCVm **260** comprises restricted file system **260A**, UCVm **262** comprises restricted file system **262A**, and UCVm **264** comprises restricted file system **264A**. Each UCVm possesses its own instance or copy of the operating system, which is isolated and separate from the main operating system (including its code and data) executing within VM0 **230** or LVM **240**. For example, UCVm **260** comprises operating system **260B**, UCVm **262** comprises operating system **262B**, and UCVm **264** comprises operating system **264B**.

To provide a low latency user experience, UCVms may not be booted from scratch each time an application is needed to be started. Instead, a UCVm may be created very quickly by cloning the UCVm from a template VM (with a booted OS) that has been pre-loaded in memory at system boot time. In an embodiment, the template used to clone a UCVm may be selected from templates **238** stored in VM0 **230**. A variety of techniques can be employed to make this cloning operation as fast as a few 100 milliseconds. Multiple types of templates may be used by a system to create UCVms depending the nature and type of application(s) to be run inside the UCVm, as discussed in greater detail below in the section entitled "Cloning a UCVm from a Template."

Cognitive assist module **236** is software that is responsible for implementing the rules and policies of embodiments as well as helping the user of client **200** in understanding and navigating the security model employed by client **200** on an as-needed basis. Cognitive assist module **236** helps decide what activities run in which UCVms, including when VMs are created or destroyed, and what kind of access to network and file system resources each UCVm has. Cognitive assist module **236** also helps protect the user, e.g., when a user is fooled by malware running in a UCVm and is in the process of providing some information that they have previously provided to enterprise code running in LVM **240** (for example a password), then cognitive assist module **236** may detect this situation and prevent the user from providing the information (which may be secret corporate information) to the malware.

Regarding the restricted file system of each UCVm, each UCVm has access to a private copy of a subset of the files in file system **242** on client **200**. A UCVm may only have access to those files which the UCVm should need for the correct operation of the application executing therein. For example, user files are usually not required for correct

12

operation of an application executing in a UCVm and thus are not typically exposed to a UCVm. On the other hand, if a UCVm is created as a result of the user wishing to edit a document using an application, such as MS Word, then a copy of the document the user wishes to edit will be provided to the restricted file system of the UCVm at the time the UCVm is created. Advantageously, using UCVm **260** as an example, if a process executing within UCVm **260** makes any changes to any files in restricted file system **260A**, then these changes do not impact the files stored in file system **242** maintained in LVM **240** because such changes are only made to restricted file system **260A** maintained in the UCVm and are not propagated, without express consent from the user, to file system **242** maintained by LVM **240**.

In a typical use case of a UCVm, the UCVm may run a local application or an individual web page session. When a user is done running the local application or navigates away from a web page to another page with a different Internet URL domain, the corresponding UCVm is destroyed. Any new local application or web application will be run inside a brand new, separate UCVm that is cloned again from a clean UCVm master template. Thus, if there has been any compromise to the UCVm during the course of running some malicious code that was introduced into the UCVm, then the adverse affects of the security breach are isolated to only the affected UCVm and are lost when the UCVm is destroyed.

For example, assume that a user double-clicks on a MS Word document icon in Windows Explorer. Embodiments create a special UCVm to run the MS Word process. In a particular embodiment, cognitive assist module **236** of VM0 **230** may dynamically create the UCVm using a template in one or more templates **238** or use a pre-existing template in memory or on the disk. The template selected by cognitive assist module **236** may be selected based on what activity is to occur within the UCVm, i.e., the selected may be designed to create a UCVm having characteristics that are optimal for running a text editor therein. The created UCVm contains a copy of the operating system as well as a restricted (local) copy of the file system. This local copy of the file system in the UCVm contains all the usual Windows and Program files; however, the user's profile folder in the local copy of the file system contains only the single target MS Word document being opened.

As another example, assume that three tabs are open in a web browser and further assume that each tab is open at a different web page. In consideration of the code which may be contained or embedded on a web page, each web page may be properly considered a web application. In embodiments of the invention, the code responsible for rendering the user interface (UI) of the web browser runs in VM0 **230**. On the other hand, executable code for the three web applications runs in three separate UCVms. A core HTML/Javascript engine runs in each of the three UCVms. A copy of the file system within each of the three separate UCVms does not contain any part of the user's files, as they are not required for the task performed by each UCVm, namely displaying a web page. Thus, each web application (or web page in this example) is completely isolated from the rest of the system.

In an embodiment, a UCVm may be connected to the Internet according to an access policy determined by the nature of the code running within the UCVm. To illustrate, web pages are typically restricted as per a strict "same origin policy" similar to the rules implemented by modern web browsers. In the "same origin policy," scripts running on

web pages are permitted to access methods and properties of other scripts originating from the same site with no specific restrictions, but are prevented from accessing most methods and properties across web pages on different sites. Untrusted native applications running outside of the web browser are restricted by default to be able to connect only to the domain from which the program was downloaded (and to specific content delivery networks (CDNs) that may be in use by the domain in question).

This level of network access for downloaded applications can be explicitly changed (increased or decreased) by the end-user to include additional sites on the Internet. End-user control over what a UCVL can connect to may be subject to certain limitations related to corporate networks and sensitive web sites (such as a bank and web mail provider). For example, any code running in a UCVL may not, in general, access any site on a corporate Intranet to which client 200 is connected. Applications that need to connect to the corporate Intranet may need to be signed by the IT administrator of the domain. Similarly, non-web untrusted application code in a general UCVL may not connect to a web site associated with a search engine or bank or other sites that may have been previously identified as being "off limits." These connections can only be made through a web browser (which spawns UCVLs bound to these special domains) or from a special purpose LVM called a VVM, which described in further detail below.

In an embodiment, there is no communication channel available for an application running in one UCVL to communicate with an application running in another UCVL. Thus, applications running in UCVLs are completely isolated from each other and from the other applications in the system. This is well suited for running downloaded third party local applications which are generally designed to be self-contained or for Internet applications (web pages are not supposed to rely on any communication between applications within the web browser). In an alternate embodiment, communication between an identified set of virtual machines can be enabled by a person with sufficient privileges, such as an IT administrator for client 200.

Firewall Virtual Machine

In an embodiment, the implementation of the network access restrictions is done in a dedicated VM called a firewall VM. FIG. 2 depicts an exemplary firewall VM 250 of an embodiment. Firewall VM 250 runs an isolated operating system with a dedicated and fixed set of firewall applications that implement the network access policy for all VMs in client 200 (except perhaps VM 230, which may not have any network access). Firewall VM 250 may provide, to any virtual machine running on client 200 in which untrusted code is executed or untrusted data is being interpreted, restricted access to only those network resources deemed necessary on an as-needed basis in accordance with a policy described by policy data stored on client 200.

In another embodiment of the invention, the firewall functionality of the system may be co-located and implemented inside either the hypervisor 220 of FIG. 2, or inside the LVM 240 of FIG. 2 (working in conjunction with the hypervisor 220 of FIG. 2), or inside VM 230 of FIG. 2 (working in conjunction with the hypervisor 220 of FIG. 2).

Validated Virtual Machines—VVMs

UCVMs are not appropriate to run local applications that interact heavily with each other using local APIs such as

COM, as typically there is no communication channel available for an application running in one UCVL to communicate with an application running in another UCVL. Embodiments may employ one (or more) special UCVLs called a Validated Virtual Machine (VVM) for the purpose of running relatively trusted local applications that have complex interactions between the applications. Such complex interactions are common in enterprise frameworks containing multiple applications, such as Microsoft's Office Suite and IBM's Lotus Notes.

FIG. 2 depicts an exemplary VVM 266 of an embodiment. Note that while FIG. 2 depicts a single VVM for ease of explanation, other embodiments of the invention may employ two or more VVMs or no VVMs based upon the particular needs of the user and/or policies of the organization responsible for or the owner of client 200.

Applications need to be signed and configured for co-location in the same VM by an administrator of client 200 before they can run in VVM 266. Inside VVM 266, signed applications can interact with each other using all types of APIs and frameworks supported by the OS being used. In an embodiment, the default network access policy of a VVM is to allow access to a corporate network only. The IT administrator may increase or decrease this level of access, subject to certain restrictions.

In an embodiment, specific signed applications or suites (groups of applications) that originate from a trusted source (other than the enterprise) may also be designated to run together in a particular VVM responsible for applications originating from that source. For example, all non-corporate applications that are signed by a specific vendor may be run together in a single VVM. These applications would then be isolated from corporate applications and general untrusted applications, but not from one another. A specific network access rule that is more permissive than the "same origin policy" used for web applications and unsigned applications may be used for a VVM. The restricted copy of file system 242 exposed to a VVM is similar to that exposed to a generic UCVL in that the restricted copy of file system 242 exposed to a VVM comprises only those files related to, or required for, performance of the applications executing within the VVM.

The Restricted File System Exposed to a VM

FIG. 3 is block diagram of the functional components involved in exposing a restricted copy of file system 242 to different UCVLs (and VVMs) according to an embodiment of the invention. File System Switch 310 is software that is configured to provide the newly created UCVL with access to a copy-on-write clone of the OS image that the UCVL was created from once the UCVL has started. The minimal operating system and program files 330 in the copy-on-write clone may be created from either the corporate LVM OS image 320 or a separate generic stripped down OS image 322 which may be created by the IT administrator.

Furthermore, a newly created UCVL is provided a copy of necessary user files 340, which are a subset of the user files in file system 242. The composition of necessary user files 340 will be different for each user. The set of files comprising the user files in file system 242 maintained in LVM 240 are typically those files in the user's home folder, e.g., c:\Users\<username>. The particular copies of files that are provided to a particular UCVL as necessary user files 340 are the minimum set of files that are needed by that UCVL to accomplish what the user intended to do as captured when the target application was being invoked. For

15

example, if the user double clicked on a specific MS Word file named ABC.docx at the location c:\Users\<username>\Documents in the file system 240 maintained in LVM 240, then necessary user files 340 would only include a copy-on-write clone of the ABC.docx file and only this copy-on-write clone of the ABC.docx file is made available in the virtual c:\Users\<username>\Documents folder made visible to the newly created UCVM running the MS Word application. If a program (like MS Word) was started without any association with a file, then necessary user files 340 would correspond to an empty c:\Users\<username>\Documents virtual folder.

Any application running in a UCVM therefore only has access to the particular set of user files provided explicitly by the user when the program was invoked. Subsequently, if the user wants to browse file system 242 for another file from within the application (for example, by using the File→Open menu item of MS Word), then he or she will see a restricted user files directory.

To enable the user to select files from the user's own User Files folder in file system 242 maintained in LVM 240 using an application executing within an UCVM, a user interface may be provided to allow the user to browse his or her files in file system 242, select one or more of the user files, and expose a copy of the selected files to the appropriate UCVM. For example, FIG. 4 is a flowchart illustrating the steps involved in a UCVM obtaining a copy of a new user file maintained in file system 242 according to an embodiment of the invention. In step 410, a special file is provided to each UCVM. The special file may be provided to the UCVM in a number of different ways, e.g., the special file may be inserted into each folder of the virtual C:\Users\<username> directory provided to each UCVM. This special file may be named something akin to "Show All My Files" or the like, as its selection will be used to trigger exposing additional copy-on-write clones of files stored in file system 242 to the UCVM.

In step 420, File System Switch 310 detects when the special file is selected by the user. For example, when a program executing within a UCVM browses to the special file, presumably as a result of a user click, this action may be trapped by File System Switch 310.

In step 430, File System Switch 310 invokes a dialog with LVM 240 that allows the user to browse the full file system 242 maintained in LVM 240. The user may then select a file or folder in file system 242. Note that at this stage, the user may be granted read access to the full file system 242 for purposes of selecting a file or folder, but the user is not granted write access to file system 242. Therefore, the user is prevented from modifying file system 242 maintained by LVM 240 in any way.

In step 440, after the user selects a file or folder, a copy of the selected file or folder is created. The copy of the selected file or folder is then inserted into the restricted file system associated with the UCVM. As a result of inserting the copy of the selected file or folder in the restricted file system associated with the UCVM, an application executing in the UCVM may have read and write access to the copy of the selected file or folder in the virtual file system, but is prevented from effecting any change to the original copy of the selected file or folder in file system 242 maintained by LVM 240.

The steps of FIG. 4 ensure that files in file system 242 maintained by LVM 240 are not visible to a UCVM without explicit permission from the user. Malicious code running in a UCVM, for example, cannot programmatically access files in file system 242 in LVM 240. Further, malicious code

16

running in a UCVM also cannot render a false user interface to trick the user into unintentionally providing any user files to the malicious code, since all code responsible for rendering the user interface is maintained within VM0 230, and thus, unreachable and un-hackable by the malicious code.

File System Switch 310 may be implemented in a variety of ways. For example, in one embodiment, File System Switch 310 may be implemented by a network file system protocol (NFS or CIFS may be used). A special VM (or LVM 240) may be used as the OS serving the "User Files" shared file system. Other VMs "mount" this shared file system using NFS or CIFS (or another network file system) from the hosting VM. Application software in the hosting VM may decide what files are exposed to which VM based on instructions provided by VM0 230.

In another embodiment, File System Switch 310 may be implemented, in part, by a proprietary protocol for handling communications between the different UCVMs and File System Switch 310. File System Switch 310, in such an embodiment, may be implemented as part of a special VM or in LVM 240.

Cloning a UCVM from a Template

In an embodiment of the invention, every virtual machine created in client 220 is instantiated using a template selected from one or more templates 238 stored in VM0 230. In an embodiment, each template in one or more templates is either immutable or may be updated in a very controlled fashion.

Each of one or more templates 238 may be used to instantiate or create a virtual machine with different characteristics or operational parameters. The characteristics or operational parameters described by a template may be configured, tailored, or suited for a particular context or type of processing activity. For example, each template may specify what type of code is to be run within a virtual machine created using the template, a size of the virtual machine created using the template, firewall settings for the virtual machine created using the template, what type of virtual machine (for example, a VVM, UCVM, or a LVM) is to be created using the template, how changes to a local file system within the virtual machine created using the template are to be persisted, and what portion, if any, of the network can a virtual machine created using the template access.

One or more devices internal to client 200 or externally connected to client 200 may interact with one or more processes executing in a virtual machine within client 200. In an embodiment, a template may assign responsibility for a selected set of devices to a virtual machine created using the template. In other embodiments, responsibility for a selected set of devices may be assigned to a particular virtual machine by virtue of policy data stored on client 200. Such policy data may describe one or more policies provided to client 200 from an owner or responsible organization of client 200. Policy data of this nature may be maintained by VM0 230 or LVM 240, for example, in certain embodiments.

In an embodiment, one or more templates 238 may be arranged in a hierarchy such that there is a root node corresponding to a template having a default set of characteristics. The root node may have one or more child nodes, and each of these child nodes may be associated with a template that inherits the properties of the parent template, but contains additional or changes properties associated with that child node. Naturally, each child node may also have

children, and so the hierarchy of templates may be an arbitrary number of levels deep, where each template inheriting characteristics of its parent, but yet each template is capable of further defining or changing characteristics that distinguishes the template over its parent.

Branches of the hierarchy of templates may be associated with, or more particularly suited, different types of activity. For example, certain templates may be associated with corporate activity, and may therefore specify characteristics related to virtual machines running corporate applications. Similarly, certain templates may be associated with the user's personal application's activity or Internet/Web related activity, and may therefore specify characteristics related to virtual machines running the user's own applications or Internet/Web applications respectively.

FIG. 5 is an illustration of instantiating a plurality of different virtual machines using different templates according to an embodiment of the invention. In FIG. 5, CVM-0 represents a template that defines a virtual machine having characteristics suitable for running a corporate application, PVM-0 represents a template that defines a virtual machine having characteristics suitable for running a user application (non-corporate), and WVM-0 represents a template that defines a virtual machine having characteristics suitable for running an Internet application. Other embodiments of the invention may define a variety of other templates to define different types of templates. In the example of FIG. 5, cognitive assist module 236 in VM0 230 may use CVM-0 to instantiate one or more corporate virtual machines, such as CVM-1, CVM-2, etc. Similarly, cognitive assist module 236 may use PVM-0 to instantiate one or more personal (non-corporate) virtual machines, such as PVM-1, PVM-2, etc., and cognitive assist module 236 may use WVM-0 to instantiate one or more web-based virtual machines, such as WVM-1, WVM-2, etc. As depicted in FIG. 5, each instantiated UCVm connects to an external network through Firewall VM 250. Cognitive assist module 236 can either create these templates on demand or create and store them while monitoring the usage of the client.

Installation of Software

In the normal operation of a typical PC, a fair amount of after-market software is installed. Such after-market software installed on a PC generally falls into one of two categories, namely (a) validated software (packages or straight executables) installed by the IT administrator of the PC or (b) end-user installed software (including web browser plugins & extensions, more complex software packages that go through an explicit install phase, and straight executables that can be executed without an explicit installation phase). Note that end-user installed software may be signed (by a verifiable, known vendor) or unsigned.

In embodiments of the invention, installation of validated software is performed as is normally performed today. The IT administrator can manage corporate validated software using embodiments using similar procedures as performed today, except that such corporate validated software are installed in LVM 240 (or if need be, VVM 266).

With respect to end-user installed software, IT administrators have two choices for how they would like to handle this type of installation using embodiments of the invention. The first choice is for the IT administrator to lock down client 200 by disallowing any installation of end-user installed software. While this is a safer operating decision, this approach may reduce the end-user's productivity because the end user cannot take advantage of applications

that may be otherwise useful that have not yet been validated by the IT administrator. The IT administrator may provide installation support on an individual and as-needed basis whenever a user wishes to install any end-user installed software; however, doing so will increase the cost of support by the IT administrator.

The second choice is for the IT administrator to allow the user to install end-user installed software him or herself using features provided by embodiments of the invention. End-user installed software may include browser plugins, browser extensions, signed and unsigned installation packages, and straight executables. Browser plugins are installed into an installed browser plugin database that is maintained in a particular UCVm. The installed browser plugin database may be implemented, in an embodiment, using file and registry diff store 820 shown in FIG. 8, which is an illustration of safely installing an untrusted application according to an embodiment of the invention of the invention. During installation of a plugin, the installed browser plugin database is also updated to record the domain that was used to initiate the plugin install. Presumably, this is the web page that contains an element or component that requires the plugin to render the complete content in the web page. Subsequently, the web browser loads an installed plugin into a web HTML/JS engine instance (which runs inside a UCVm) only if the domain of the web page to be displayed by the UCVm matches a domain, recorded in the plugin database, associated with the installed plugin. A plugin that is used by multiple sites is installed only once, but is associated with multiple domains. Popular plugins like Flash may be pre-installed in certain embodiments.

Browser extensions may be installed into a web browser's extension database that is maintained in a particular UCVm. During runtime, browser extensions are treated like web applications in that each browser extension is run inside its own UCVm. In an embodiment, the web browser extension database and the installed browser plugin database may be implemented in the same database in a single UCVm.

Signed installation packages may be run and the resulting installation may update either the LVM image or the Generic Windows image based on a policy set by the IT administrator.

Unsigned installation packages go through a virtual install. The virtual installation of unsigned installation packages will be described with reference to FIG. 8, which is an illustration of safely installing an untrusted application according to an embodiment of the invention of the invention. A registry and program files change set is created and stored in file and registry diff store 820. Start-menu and desktop changes by the installer are captured in a special folder which contains desktop and start menu items for all user-installed applications. Subsequently, if an unsigned application is run, it is run in a UCVm cloned from the Generic Windows image all by itself. Virtual disk 810 in FIG. 8 is the normal virtual disk of the UCVm. DiffStore 820, which further virtualizes the file system and the registry as seen by the applications of UCVm, is typically implemented as a separate module outside of the normal block level virtual disk store.

Signed and unsigned executables may be run in a UCVm. Such a UCVm may be created on demand and destroyed after its use is ended by embodiments.

Managing Web Cookies and Caches

A web cookie (or simply "cookie") is a piece of text stored on a user's computer by their web browser. A cookie can be

used for authentication, storing web site preferences, shopping cart contents, the identifier for a server-based session, or anything else that can be accomplished through storing text data.

While the actual cookie itself is not visible to the user, the user would notice a difference in the user experience of interacting with a web site if cookies could not be saved between visits to the web site. Accordingly, embodiments of the invention provide mechanism to store cookies before a UCVm is destroyed, so that the next time the user visits the web site using a web browser running in a different UCVm, any cookies that have been stored and are associated with that web site may be injected into the new UCVm.

Similarly, to provide the best user experience, it would be advantageous to carry over the cache of a web browser for a particular web domain from one UCVm to the next, so that the next time the user visits the web domain using a different UCVm, there is no a delay in displaying content due to an unpopulated cache. Thus, embodiments of the invention provide mechanism to store the web cache of a web browser for a web domain before a UCVm is destroyed, so that the next time the user visits the web site using a web browser running in a different UCVm, the cache of the web browser need not be warmed (i.e., repopulated), as the cache in the new UCVm has been updated to contain all the objects the cache previously contained in the prior, and now destroyed, UCVm used to visit the web domain.

To provide a concrete example with reference to the example of FIG. 2, assume that a user initially transparently uses UCVm 260 to run a web browser to visit web site A. When UCVm 260 is destroyed, any cookies and cache files are extracted and saved. Thereafter, assume the user transparently uses UCVm 262 to run a web browser to visit web site B. As web site B is hosted by a different web domain than web site A, the previously stored cookies and cache files associated with web site A will not be injected into UCVm 262. Thereafter, if UCVm 262 is destroyed, then any cookies and cache files are extracted and saved. At a later point in time, if the user thereafter transparently uses UCVm 264 to run a web browser to visit web site A, then the previously stored cookies and cache files associated with the web domain of web site A will be injected into UCVm 264. This allows the web browser running in UCVm 264 to visit web site A to appear, to the user, to have the same state of the prior web browser used to visit web site A, even through different virtual machines are used between visits. Note that no portions of the file system are saved between visits to a web site; only the state of the web session is saved.

In one embodiment, the cookies and cache information is captured in Diff Store 820 associated with the URL of the website. In each visit to the same URL, the UCVm utilizes the same Diff Store presenting the cookies and caches to the UCVm. In another embodiment, the cookies and cache files can be captured at the end of the session and saved to the client system's core file system in a special folder. On visiting the same URL again, the cookies and cache can be re-injected into the file system of the UCVm.

Efficient Physical-to-Virtual Disk Conversion

Platform virtualization is performed on a given hardware platform by host software (a control program), which creates a simulated computer environment, a virtual machine, for its guest software. A hypervisor, also called virtual machine manager (VMM), is one of many hardware virtualization techniques that allow multiple operating systems, termed guests, to run concurrently on a host computer. The hyper-

visor presents to the guest operating systems a virtual operating platform and manages the execution of the guest operating systems. A guest OS executes as if it was running directly on the physical hardware. Access to physical system resources such as the network access, display, keyboard, and disk storage is suitably virtualized so that guest OS does not know these are virtual devices.

Generally, there are two types of hypervisors. Type 1 (or native, bare metal) hypervisors run directly on the host's hardware to control the hardware and to manage guest operating systems. A guest operating system thus runs on another level above the hypervisor. Type 2 (or hosted) hypervisors run within a conventional operating system environment. With the hypervisor layer as a distinct second software level, guest operating systems run at the third level above the hardware. In other words, Type 1 hypervisor runs directly on the hardware; a Type 2 hypervisor runs on another operating system, such as Windows. Embodiments of the invention may use any type of hypervisor. Thus, hypervisor 220 in FIG. 2 may either be a Type 1 or a Type 2 hypervisor.

A virtual disk image is a file on a physical disk, which has a well-defined (published or proprietary) format and is interpreted by a hypervisor as a hard disk. In terms of naming, a virtual disk image may have a specific file type extension, e.g., .vmdk for VMware VMDK, .vhd for Xen and Microsoft Hyper-V, and .vdi for Oracle VM VirtualBox.

There are two approaches employed for storage allocation by prior hypervisors, namely, (1) pre-allocate the entire storage for the virtual disk upon creation and (2) dynamically grow the storage on demand. In the former approach involving pre-allocation, the virtual disk may be implemented as either split over a collection of flat files (typically one is 2 GB in size) or as a single, large monolithic flat file. In the latter approach involving on-demand growth, the virtual disk may also be implemented using split or monolithic files, except that storage is allocated on demand.

There are two modes in which a disk can be mapped for use by a virtual machine. In a virtual mode, the mapped disk is presented as if it is a logical volume, or a virtual disk file, to the guest operating system and its real hardware characteristics are hidden. In a physical mode, also called the pass through mode, the hypervisor bypasses the I/O virtualization layer and passes all I/O commands directly to the disk.

A virtual machine (VM) is a software implementation of a machine (i.e. a computer) that executes programs like a physical machine. Virtual machines allow the sharing of the underlying physical machine resources between different virtual machines, each running its own operating system. The software layer providing the virtualization is called a hypervisor, such as hypervisor 220 in FIG. 2.

Virtual machines each require their own image of the operating system. The guest OS and host OS typically do not share the OS image, even if they are the same OS. This is problematic for several reasons. First, if a user wishes to run 10 different virtual machines, then she will require 10 different copies of the OS for the guest OSs, which requires an undesirable amount of storage to maintain. As she is already running one virtual machine at the host, the total number of different copies of the OS required would be 11. Second, the OS for a VM has to be created either by installing a new OS or shipping a copy of the OS from somewhere else, which is burdensome for those who do not have access to OS images. Further, it is also time consuming to install a new OS or ship an OS image, which is typically quite large. A third problem is that any software present in

the host OS (such as a printer driver) will not be available in a guest OS unless it is installed again.

Shadow Copy (Volume Snapshot Service or Volume Shadow Copy Service or VSS) is a technology included in Microsoft Windows that allows taking manual or automatic backup copies or snapshots of data (termed “shadow copies”), even if it has a lock, on a specific volume at a specific point in time over regular intervals. VSS operates at the block level of the file system. Shadow Copy is implemented as a Windows service called the Volume Shadow Copy service. Software VSS provider service is also included as part of the Microsoft Windows OS to be used by Windows applications. Shadow Copy technology requires that the file system to be NTFS to be able to create and store shadow copies. Shadow Copies can be created on local and external (removable or network) volumes by any Windows component that uses this technology, such as when creating a scheduled Windows Backup or automatic System Restore point.

Snapshots have two primary purposes. First, they allow the creation of consistent backups of a volume, ensuring that the contents cannot change while the backup is being made. Second, they avoid problems with file locking. By creating a read-only copy of the volume, backup programs are able to access every file without interfering with other programs writing to those same files. Through the integration between the Volume Shadow Copy Service, hardware or software VSS providers, application level writers and backup applications, VSS enables integral backups that are point in time and application level consistent without the backup tool having knowledge about the internals of each application. The end result is similar to a versioning file system, allowing any file to be retrieved as it existed at the time any of the snapshots was made. Unlike a true versioning file system, however, users cannot trigger the creation of new versions of an individual file, only the entire volume.

Embodiments of the invention overcome this limitation by creating virtual disks based on VSS shadow copies. FIG. 6 is an illustration of a virtual disk based on VSS shadow copies according to an embodiment. The virtual disk of FIG. 6 allows for many guest OSs running on the same host to share the same OS copy with the host OS. VSS shadow copies may be created fast and efficiently. Creating virtual disks on top of VSS is also a very fast operation, which means that VMs (with same OS as host OS) can be created very efficiently. Shadow copies are also maintained cheaply by windows OS by keeping the changes since the time shadow was created. Hence, the disk usage of multiple VMs is reduced substantially. VMs can also be maintained very efficiently since VSS snapshots can be updated once and have the changes reflected in all VMs. Since a VSS shadow copy contains all the software the user has installed on the machine at the time of the VSS shadow copy creation, virtual disks also receive access to all the software. Moreover, the version of the software, including any patches installed, is exactly the same. In addition to all the software, user documents are also visible to virtual machines. A virtual disk of an embodiment is an accurate point-in-time copy of host physical disk.

In an embodiment where VSS snapshots are read-only, a ‘Delta Store Disk’ may be attached to the virtual disk. The Delta Store disk is used to capture all the changes being made to the virtual disk.

Security Afforded by Embodiments

Embodiments of the invention provide a secure environment to prevent malicious code from affecting any lasting

change in a computer system. Arbitrary code (either a web application or a native executable) runs inside an isolated operating system running on an isolated virtual machine. This code has no access to any other application (either a native application or a web application) being run by the user because those applications run in other operating systems running in separate virtual machines. Moreover, arbitrary code has access to only the specific part of the file system that is needed for correct execution of the code. Access to additional parts of the file system has to be provided by code that runs in VM0 (which is secure and fortified against unauthorized intrusion) and any increased access needs explicit authorization from the human user.

Specific trusted code that needs to interact in a complex way with other applications may be explicitly designated to run together inside the same designated VM. This type of VM also has limited access to the file system.

All code has limited network access to just what that code needs for its correct execution. All virtual machines are created from templates stored in VM0 which are either immutable or can be updated in a very controlled fashion. Consequently, if a security bug exists in a piece of code, the effect of the security bug is isolated (“space limited”) because the compromised code has access to only a limited part of the file system, the network, devices, etc. Moreover, the effect of the security bug is “time limited” because the virtual machine that has been compromised will be subsequently discarded and a new virtual machine is created for future application instances from a clean immutable VM template.

Using Policy Data to Manage the Deployment of Virtual Machines

Embodiments allow code that originates from arbitrary external sources to be safely executed by a client. In this way, digital content of unknown trustworthiness may be safely received and potentially executed and/or interpreted by a client without incurring the risk that the digital content contains malicious code that could cause undesirable consequences.

The ‘digital content’ received by the client from an external source may correspond to any type of digital data, such as executable code or non-executable, interpreted data for example. Since malicious code may be carried within certain types of non-executable data and subsequently spread when the data is interpreted by applications, embodiments treat all incoming digital content as being capable of containing malicious code, even if the digital content is not in a recognized executable form. Non-limiting, illustrative examples of digital content include an “.exe” file, an application, a collection of applications designed to run together, a portion of an application, an email attachment, a slide presentation, a text document, and a web page (which essentially is a portion of an application, namely a web browser). Even though the email attachment, the slide presentation, and the text document, in and of themselves, are not executable files, embodiments of the invention treat these forms of digital content as potentially carrying malicious code.

To manage the risk posed by receiving digital content of unknown trustworthiness, any digital content received by a client is stored in one or more virtual machines. In an embodiment, digital content received from an external source may immediately be stored in one or more virtual machines upon receipt. Alternately, digital content received

from an external source may be stored in an intermediate location, such as a local cache, prior to storing the digital content in a virtual machine.

While embodiments are configured to process all digital content originating from an external source in a virtual machine, the complexity of determining in which virtual machine the digital content should be stored and how that virtual machine should be configured is hidden from the user whenever possible or appropriate. To accomplish this goal, techniques are discussed herein for programmatically managing a plurality of virtual machines on the client to accommodate the wide variety of use cases for receiving digital content at a client. However, in some cases, explained in more detail below, it may be appropriate to inform the user of certain activity concerning a virtual machine, such as when obtaining express permission from the user is advisable before performing an action.

Certain sources of digital content are more trustworthy than other sources. For example, the web site of a bank or Fortune 500 company may be more trustworthy than the web site of a smaller company or lesser known organization. Also, applications may have different operating needs, e.g., certain applications may be designed to work closely with other applications or require access to network resources. Thus, in an embodiment, the attributes of each virtual machine are specifically tailored to reflect the type of digital content and/or applications operating or stored therein.

To illustrate how one embodiment operates, when a client determines that digital content, originating from an external source, is to be received or processed by the client, the client may identify, without human intervention, one or more virtual machines, executing or to be executed on the client, into which the digital content is to be received. To do so, the client may consult policy data, such as policy data 239 stored at client 200 of FIG. 2, to determine a placement policy, a containment policy, and a persistence policy used in identifying the one or more virtual machines into which the digital content is to be received.

The policy data may be used to specifically tailor the operation of each virtual machine to reflect the type of digital content and/or applications operating or stored therein. The placement policy identifies a particular virtual machine into which the digital content is to be stored, the containment policy identifies what network resources and client resources the particular virtual machine can access, and the persistence policy identifies whether data (or a part of it) stored in the particular virtual machine is persistently stored. Naturally, the placement policy, containment policy, and persistence policy are, to a certain extent, intertwined, as the resources a virtual machine may access and whether data stored therein is persisted will affect what applications/digital content are appropriate to reside therein.

In an embodiment, each of the placement policy, the containment policy, and the persistence policy may consider a variety of different factors. For example, the placement policy, the containment policy, and/or the persistence policy may consider a historical record of use for the client in identifying a virtual machine. The evaluation of a policy may involve consulting a historical record of how the client, or applications running thereon, has been used. In this way, if a particular action has been judged to be more safe (or less safe) over a period of time, the manner in which the action is handled by the policy may evolve over time. To illustrate, in an embodiment, if a particular network resource, such as an affiliate corporate web page, is demonstrated to be sufficiently safe over a period of time, then this web page may be processed using relaxed restrictions, e.g., by a web

browser in a virtual machine already handling another trusted web page as opposed to instantiating a new virtual machine to handle the affiliate corporate web page. On the other hand, if the historical record of use demonstrates that an action involving a particular network resource or client resource may pose some risk to the client, then the policy may subsequently handle this action more sensitively than before, e.g., by assigning code to handle the particular network resource or client resource in a dedicated virtual machine with restricted access to client and network resources.

As another example of the types of factors which may be considered by a policy, one or more of the placement policy, the containment policy, and the persistence policy may consider a current physical location of the client or to which networks the client currently has access in identifying one or more virtual machines which should be used to receive content. In this way, which networks are available to the client, the IP address assigned to the client, the current location of the client based on global positioning service (GPS) data, and the current location of the client based on an IP address or which networks are available to the client may all be considered when determining which virtual machine should receive digital content and what restrictions should be placed on that virtual machine. In this way, when the client is physically located in an area deemed safe (such as a work office or home), digital content received by the client may be handled by a virtual machine having a set of lesser restrictions than when the client is physically located in an unknown area.

As another example of the types of factors which may be considered by a policy, one or more of the placement policy, the containment policy, and the persistence policy may consider the proximity of the client to a wireless device, such as a Bluetooth enabled cell phone. For example, if the client is not within a configurable distance to the cell phone of the user of the client, then the client may receive digital content using a set of greater restrictions, e.g., code executing in all virtual machines may be denied access to certain client resources and/or all network resources. Embodiments may determine whether the client is within a configurable distance to a wireless device using a variety of different methods, such as accessing the wireless signal strength between the client and the wireless device.

In an embodiment, at least a portion of the policy data, used in identifying one or more responsible virtual machines to receive digital content, is obtained from a remote server after the client determines that digital content is to be received from an external source. In this way, policy data may be sent, as needed, from an IT administrator to the client. The client may treat any policy data already residing on the client in the same manner as policy data retrieved from a remote server. For example, when a user of the client performs an action, the client may consult a remote server to see if the remote server has any additional policy data regarding this action. Following this procedure, an IT administrator can maintain a high level of control on how the client will manage virtual machines running on the client. This enables the IT administrator to make adjustments to the security model followed by the client in real-time. The client may interact with a human operator at a remote location to obtain additional policy data or may interact with a remote automated system, without human intervention, to obtain the additional policy data. Note that certain embodiments may be configured to consult a remote server for policy data only when a certain configurable action is taken. Therefore, in certain embodiments, the client need not always contact a

25

remote server to determine if additional policy data is available each time that the client is to receive new digital content.

In an embodiment, the policy data may specify that the virtual machine assigned to receive digital content can only access a limited subset of the metadata properties for a client resource or a network resource. For example, a virtual machine may not be capable of determining what local wireless networks are available in the vicinity or whether the network card of the client is of a particular type. In this way, the amount and type of information exposed to a particular virtual machine may be controlled to a fine level of granularity.

Use of the placement policy, the containment policy, and the persistence policy by certain embodiments will be discussed in further detail below.

Placement Policy

The placement policy identifies a particular virtual machine into which the digital content is to be stored. The particular virtual machine identified by a placement policy in which digital content is to be stored may be an existing virtual machine or a new virtual machine that has not yet been instantiated. In the case where the placement policy specifies that the digital content should be received by a virtual machine that has not yet been instantiated, either the placement policy itself or some other location in the policy data will identify a template for use in instantiating the particular virtual machine. The identified template will describe characteristics of a virtual machine suitable for receiving the digital content.

The placement policy may weigh a variety of different considerations in determining which virtual machine should store the digital content so that the digital content may be safely executed, interpreted, and/or processed. For example, a placement policy of an embodiment may assign any file having a certain name or certain attributes to a virtual machine having certain characteristics. To illustrate, a placement policy may indicate that all signed executable files from an internal organization or company are to be assigned to a virtual machine having a specified set of characteristics. As another example, the placement policy may instruct untrusted applications to execute in separate virtual machines so that each untrusted application is isolated from other applications and data of the client.

The placement policy of an embodiment may identify a plurality of classes of virtual machines, where each class of the plurality of classes is associated with a different trust level for external sources of digital content. Code executing in a virtual machine cannot access external sources associated with less trustworthy external sources of digital content. For example, assume there are three classes of virtual machines, where the first class of virtual machines is designed to run web browsers accessing web sites of financial institutions and email providers, the second class of virtual machines is designed to run web browsers accessing web sites of Fortune 500 companies, and the third class of virtual machines is designed to run web browsers accessing all other web sites. In this example, a web browser executing in a virtual machine that is associated with the third class cannot access any web sites from Fortune 500 companies or financial institutions and email providers. Similarly, in this example, a web browser executing in a virtual machine that is associated with the second class cannot access any web sites from financial institutions and email providers.

26

The placement policy of an embodiment may identify the particular virtual machine into which the digital content is to be received by observing application dependencies. Such a policy recognizes that in some instances, it is helpful or even necessary to execute certain applications within a single virtual machine. For example, certain providers of software applications may design their software applications to work together or integrate with each other to a high degree. In this case, it would be advantageous to have applications that are designed to work together to run within a single virtual machine. One way for the placement policy to make this determination would be to ask the user whether an application being installed is dependent upon another application already installed at the client to ensure that both applications may be run in the same virtual machine. While this does expose the notion of a virtual machine to the user, a user need only make a decision of this nature when an application is installed on the client, and thus, this decision may be made by IT administrators or other knowledgeable personal rather than relying upon the end user of the client to make such a decision.

Alternatively, determining whether an application being installed is dependent upon another application may be made programmatically by examining the dependencies during the installation of that application. For example, during the installation of application A, the install process may check if module B is already installed or may require that module B already be installed. In this example, the placement policy may determine then that application A has a dependency with module B and may therefore allow application A to run in same virtual machine as module B.

To illustrate another example, it is initially noted that there need not be a one to one correspondence between a web browser and a web page. For example, a web browser may comprise many tabs, and each tab may display a different web page. In addition, each web browser may have a variety of different plug-in and/or associated programs which may be treated as or considered a separate application. Since a web browser may display multiple web pages of varying trust levels, it is desirable to accommodate a web browser having multiple tabs without requiring that the web pages displayed by each tab reside in the same virtual machine. For example, if a web page contains malicious code, then it would be beneficial to execute it in a different virtual machine from the virtual machine containing the web page of your bank. Therefore, in an embodiment, the placement policy may specify that web page of certain sources should be received in a separate virtual machine. While the user may see a single web browser having two tabs, on the back end this may be implemented in two separate virtual machines that each execute a copy of the web browser and possess one web page to be shown in associated with one tab of the web browser. A practical implementation of web page placement may use a VM per web-site placement policy.

These are merely examples of how a placement policy may be implemented. It is contemplated that actual implementations of a placement policy will be configured based upon the particular needs and concerns of the end user. The containment policy of certain embodiments will now be presented in greater detail.

Containment Policy

The containment policy identifies what network resources and client resources a particular virtual machine can access. Network resources, as broadly used herein, refers to any resource that is external to the client while client resources,

as broadly used herein, refers to any resources that is internal to the client. A client resource may include any device, component, and/or data residing on or accessible to the client, such as a digital camera, a network interface card, a digital clock, the current time, files, pictures, and email.

The containment policy is used to ensure that code running within a virtual machine has access to only those resources deemed necessary for normal and intended operation. For example, email attachments should not need access to the Internet (generally speaking), and so they should be opened in a virtual machine that is configured such that it does not have access to the Internet.

In an embodiment, the containment policy may specify what portion of the network that is available or exposed to code executing within a virtual machine. For example, the containment policy may specify that code executing within a particular virtual machine may access no network resources, all network resources, or a subset of the network resources. Thus, a containment policy may specify that code executing within a virtual machine may access a first set of network resources and may not access a second set of network resources. Embodiments may specify what particular network resources are available to a virtual machine using any level of granularity, e.g., only certain types of network resources may be exposed, only certain properties of network resources may be exposed, or only certain portions of the network may be exposed.

In an embodiment, enterprise applications may be grouped into collections. Groupings may be based on a variety of factors, such as job functions or business unit, for example. Each grouping of applications may be executed within a single virtual machine according to an embodiment.

To illustrate the interaction between the containment policy and client resources, the containment policy of an embodiment identifies each client resource accessible to a virtual machine. For example, a containment policy may specify whether code executing in the particular virtual machine can perform one or more of the following actions: access a USB port on the client, perform a copy operation or a paste operation, access a network to which the client is connected, access a GPS device of the client, location information for the client, or tilt information for the client, access a printer or facsimile machine to which the client is connected, and access a digital camera or screen data for the client. Note that these exemplary actions are not meant to provide an exhaustive list, as a containment policy may be used to specify, with particular specificity, which client and network resources may be accessed by code executing within a virtual machine. In this way, if a new client resource becomes available, such as fingerprint scanning device, the containment policy may be updated to reflect the new client resource available to the client.

In an embodiment involving the receipt of executable code at a client, the containment policy may specify that the executable code is denied access to a user file without first obtaining a user's permission to allow the executable code to access the user file. In this way, virtual machines may be configured to allow request permission each time executable code therein access a user file, thereby allowing the user to be informed of the intentions of the executing code and presumably prevent unauthorized access to the user's own files. Such a permission scheme might be implemented naturally as part of the normal user work flow of picking a file to open by running the permission code in a clean protected VM separate from the VM running the untrusted code which is making the request.

To illustrate the interaction between the containment policy and network resources, the containment policy of an embodiment identifies whether code executing in a particular virtual machine can one or more networks accessible to the client. As another example, the containment policy of an embodiment identifies which, if any, objects stored over a network the virtual machine can access. For example, a virtual machine may be restricted to access a specified set of objects or files on a particular server or a particular set of web pages.

In an embodiment, the containment policy may consider any number of factors, including but not limited an identity of the user of the client, a set of properties of the digital content, a physical location of the client, the current time, a holiday schedule, and a set of administrator-specified policy rules. In this way, the containment policy may assign a virtual machine having more restrictions than usual to receive digital content when the digital content is deemed more likely to contain malicious code. For example, it may be deemed likely that digital content contains malicious code when it is received by the client outside of normal business hours, over a holiday, at a time when the client is outside of the user's home or work office, or when the digital content has certain suspicious properties. In this way, the containment policy may assign suspicious digital content to be received in a virtual machine having additional restrictions appropriate for such suspicious digital content.

These examples of how a containment policy may operate and merely illustrative of some examples and are not intended to be an exhaustive list, as actual implementations of a containment policy will be configured based upon the particular needs and concerns of the end user. The persistence policy of certain embodiments will now be presented in greater detail.

Persistence Policy

In an embodiment, the persistence policy identifies whether data stored in a particular virtual machine is persistently stored. The policy grapples with the issue of whether or not to save state created by untrusted code and if so, whether the state should be stored in an isolated manner or merged back into the main file system of the computer. On one hand, to provide a convenient user experience, it may be helpful to persistently store cookies for a web site. On the other hand, it would not be desirable to persistent malicious code, such as a key logger, that was inadvertently introduced into a virtual machine by malware downloaded into and run in the affected virtual machine.

The persistence policy, hand in hand with the placement policy, should be designed to ensure that any potentially malicious code is not persistently stored, or in the alternative, persistently stored in an isolated way. This way, if malicious code, such as a key logger, is persistently stored, and in any future invocation (execution or interpretation), it is invoked (executed) in the context of a possibly new virtual machine instance separate from any other code, thereby nullifying the risk presented thereby.

To illustrate an illustrative persistence policy, in an embodiment only cookies and cache files are persistently stored in a virtual machine in which a web browser executes. Further, the cookies and cache files associated with a particular web site are only inserted to a virtual machine that is intended to execute a web browser displaying that web site. Thus, cookies and a cache file associated with site A would not be inserted into a virtual machine instantiated to run a

29

web browser to display web site B, but would be inserted into a virtual machine instantiated to run a web browser to display web site A.

The above discussion of a persistence policy is exemplary of certain embodiments and is not intended to describe all implementations of a persistence policy, as a persistence policy will be configured based upon the particular needs and concerns of the end user.

Unified Display

Even though there may be a plurality of virtual machines executing at the client, this complexity need not be exposed to the end user of the client. Thus, the end user should be presented visual content generated from each virtual machine executing on the client in a unified manner to present a single, cohesive presentation to the end user of the client. The presentation of the content should be seamless and close to native as possible.

For example, the end user of the client should interact with a web browser that looks like a known web browser, even though the web browser, at the back end, is implemented using a plurality of virtual machines to execute copies of the web browser and different web pages corresponding to each tab of the web browser.

Hardware Mechanisms

In an embodiment, client 200 of FIG. 2 may be implemented on, include, or correspond to a computer system. FIG. 9 is a block diagram that illustrates a computer system 900 upon which an embodiment of the invention may be implemented. In an embodiment, computer system 900 includes processor 904, main memory 906, ROM 908, storage device 910, and communication interface 918. Computer system 900 includes at least one processor 904 for processing information. Computer system 900 also includes a main memory 906, such as a random access memory (RAM) or other dynamic storage device, for storing information and instructions to be executed by processor 904. Main memory 906 also may be used for storing temporary variables or other intermediate information during execution of instructions to be executed by processor 904. Computer system 900 further includes a read only memory (ROM) 908 or other static storage device for storing static information and instructions for processor 904. A storage device 910, such as a magnetic disk or optical disk, is provided for storing information and instructions.

Computer system 900 may be coupled to a display 912, such as a cathode ray tube (CRT), a LCD monitor, and a television set, for displaying information to a user. An input device 914, including alphanumeric and other keys, is coupled to computer system 900 for communicating information and command selections to processor 904. Other non-limiting, illustrative examples of input device 914 include a mouse, a trackball, or cursor direction keys for communicating direction information and command selections to processor 904 and for controlling cursor movement on display 912. While only one input device 914 is depicted in FIG. 9, embodiments of the invention may include any number of input devices 914 coupled to computer system 900.

Embodiments of the invention are related to the use of computer system 900 for implementing the techniques described herein. According to one embodiment of the invention, those techniques are performed by computer system 900 in response to processor 904 executing one or

30

more sequences of one or more instructions contained in main memory 906. Such instructions may be read into main memory 906 from another machine-readable medium, such as storage device 910. Execution of the sequences of instructions contained in main memory 906 causes processor 904 to perform the process steps described herein. In alternative embodiments, hard-wired circuitry may be used in place of or in combination with software instructions to implement embodiments of the invention. Thus, embodiments of the invention are not limited to any specific combination of hardware circuitry and software.

The term "machine-readable storage medium" as used herein refers to any tangible medium that participates in storing instructions which may be provided to processor 904 for execution. Such a medium may take many forms, including but not limited to, non-volatile media and volatile media. Non-volatile media includes, for example, optical or magnetic disks, such as storage device 910. Volatile media includes dynamic memory, such as main memory 906.

Non-limiting, illustrative examples of machine-readable media include, for example, a floppy disk, a flexible disk, hard disk, magnetic tape, or any other magnetic medium, a CD-ROM, any other optical medium, a RAM, a PROM, and EPROM, a FLASH-EPROM, any other memory chip or cartridge, or any other medium from which a computer can read.

Various forms of machine readable media may be involved in carrying one or more sequences of one or more instructions to processor 904 for execution. For example, the instructions may initially be carried on a magnetic disk of a remote computer. The remote computer can load the instructions into its dynamic memory and send the instructions over a network link 920 to computer system 900.

Communication interface 918 provides a two-way data communication coupling to a network link 920 that is connected to a local network. For example, communication interface 918 may be an integrated services digital network (ISDN) card or a modem to provide a data communication connection to a corresponding type of telephone line. As another example, communication interface 918 may be a local area network (LAN) card to provide a data communication connection to a compatible LAN. Wireless links may also be implemented. In any such implementation, communication interface 918 sends and receives electrical, electromagnetic or optical signals that carry digital data streams representing various types of information.

Network link 920 typically provides data communication through one or more networks to other data devices. For example, network link 920 may provide a connection through a local network to a host computer or to data equipment operated by an Internet Service Provider (ISP).

Computer system 900 can send messages and receive data, including program code, through the network(s), network link 920 and communication interface 918. For example, a server might transmit a requested code for an application program through the Internet, a local ISP, a local network, subsequently to communication interface 918. The received code may be executed by processor 904 as it is received, and/or stored in storage device 910, or other non-volatile storage for later execution.

Clipboard Hardening Across Virtual Machines

In an embodiment, OS 232 may implement or support a clipboard that may be accessed by multiple virtual machines of client 200, such as UCVMs 260, 262, and 264. To illustrate a concrete example, consider that the Microsoft

31

Windows operating system supports a Windows clipboard. OS 232 may correspond to the Microsoft Windows operating system. Objects may be copied to the clipboard of OS 232. Once an object is copied to the clipboard OS 232, it may be copied (or “pasted”) to another location, such as the clipboard of the operating system running in UCVMs 260, 262, and 264. In this way, the clipboard of OS 232 may be used to copy objects from one location to another, and thus, potentially from one virtual machine to another or between a virtual machine and the host.

Copying objects to and from the clipboard of OS 232 may be performed transparently to the user (and thus without human intervention or instruction), but only if deemed permissible by a policy. When a user causes an object to be copied into the clipboard of a guest operating system, a software module executing on an operating system of client 200 may access a policy to determine if the object should be copied to another clipboard residing elsewhere. For example, a policy may be used to determine whether an object copied into the clipboard of a guest operating system should be copied into the clipboard of a host operating system. The clipboard of the host operating system may be accessible to a variety of different processes. Objects in the clipboard of a host operating system may be caused to be copied into the clipboard of another guest operating system running in a virtual machine.

In an alternate embodiment, rather than using a policy to regulate which objects are copied into a clipboard residing in the host operating system, a policy may be used to regulate which objects are copied into a clipboard residing in a guest operating system running in a secure virtual machine. Such a virtual machine may be designed to function as a global clipboard for virtual machines and the host operating system running on client 200.

Embodiments define policies in policy data 239. These policies are enforced by software module 236. The policies defined in policy data 239 ensure that objects copied into the clipboard from one virtual machine (or the host) cannot be copied into another virtual machine (or the host) in an inappropriate manner. Such embodiments advantageously prevent potential security breaches. To illustrate, if one virtual machine is compromised, then malicious code in that virtual machine code may steal a bank password or the like copied to the clipboard by another virtual machine or the host.

However, there are legitimate reasons for wanting to share objects in the clipboard between virtual machines in some cases. Certain applications may require this functionality for normal operation. Also, if a human operator wishes to copy objects from the clipboard, then the operation should be allowed since this action is not requested by malware. Accordingly, in certain embodiments, policy data 239 may define a policy that distinguishes between human interaction and automated clipboard access. If it is determined that a human operator is initiating an action involving the clipboard at client 200, then objects in the clipboard may be shared between virtual machines and/or the host. On the other hand, if code executing on client 200 is determined to be requesting an object in the clipboard that was copied there from another virtual machine, then if the policy dictates so, the request will be denied to prevent the code from accessing the object in the clipboard under the rationale that the code may be compromised by or may correspond to malware. Similarly, the policy of an embodiment may be defined such that the policy disallows copying of clipboard content to a virtual machine if this clipboard was populated by code running on the operating system of client 200, even in case

32

of explicit human request. Such a policy allows copying of clipboard content between virtual machines, but prevents a human operator from mistakenly sharing the clipboard of client 200 (that might be confidential in nature) with a malicious virtual machine.

In an embodiment, input received from a physical mouse or the keystrokes made on the physical keyboard will be treated as originating from a human user. Embodiments attempt to distinguish between instructions submitted via actual physical, tangible input mechanisms and spoofed instructions that masquerade as legitimate input submitted by malicious code. As an example, the keystrokes and clicks of the mouse may be analyzed to determine whether a human user has issued them with the intention of performing a clipboard operation. Policy data 239 may define various conditions or tests for determining whether a request to access a clipboard originated from a human user.

Embodiments of the invention need not be limited to the use of a clipboard, but instead, these techniques may be broadly applied to any type of shared memory. Embodiments may use policies to control the access to shared memory from a plurality of virtual machines based upon whether the access is perceived to originate from a human user or executing code.

Embodiments of the invention acknowledge that analyzing mouse click semantics in an opaque guest desktop may be onerous; consequently, consequently, certain embodiments may perform certain menu processing functionality within the host, where such reasoning and analysis is less onerous and reliable.

On certain operating systems, such as Microsoft Windows, a single logical clipboard operation may involve multiple physical data transfers, and accordingly, these separate data transfers need to be recognized as belonging to a single clipboard operation and treated as such.

On certain operating systems, such as Microsoft Windows, a clipboard can simultaneously hold multiple objects that belong to different clipboard object types. Some of these object types (such as, for example, plain text or an image bitmap) are simple and have well defined semantics, while other object types are complex and their semantics may not be clearly documented. Exposing trusted code to objects of unknown semantics may potentially result in a breach of security of the system. Therefore, if a policy dictates so, certain types of clipboard objects are not allowed to be transferred from one clipboard to another even if the object is copied into the original clipboard as the request of a human operator. In this way, a policy may prevent a user from inadvertently introducing a potentially hazardous type of object into a clipboard accessible to sensitive parts of client 200 (such as OS 232), even if the user explicitly makes a request to do so. In effect, such a policy allows only object types deemed “safe” to be copied from potentially unsafe sources (such as a virtual machine designed to run untrusted code or interpret untrusted data) to another location (such as the clipboard of OS 232) accessible to trusted portions of client 200.

Embodiments of the invention require some human interaction to either initiate or approve copying objects from one clipboard to another clipboard. An example of a human interaction sufficient to initiate copying might be a key stroke and an example of a human interaction sufficient to approve copying might be express consent of a user submitting via a host dialogue interface. Similarly, in an embodiment, objects are only copied to a guest clipboard

33

(from either the host operating system or another virtual machine) if the user has either initiated or approved a paste operation.

In the foregoing specification, embodiments of the invention have been described with reference to numerous specific details that may vary from implementation to implementation. Thus, the sole and exclusive indicator of what is the invention, and is intended by the applicants to be the invention, is the set of claims that issue from this application, in the specific form in which such claims issue, including any subsequent correction. Any definitions expressly set forth herein for terms contained in such claims shall govern the meaning of such terms as used in the claims. Hence, no limitation, element, property, feature, advantage or attribute that is not expressly recited in a claim should limit the scope of such claim in any way. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

What is claimed is:

1. One or more non-transitory machine-readable storage mediums storing one or more sequences of instructions, which when executed by one or more processors, causes:

upon a software module, executing on a first operating system, determining that the first operating system has caused data to be written to a first clipboard maintained by the first operating system, the software module consulting policy data to determine whether, without ceasing execution and without changing priority of one or more of said first operating system and a second operating system, the data should additionally be written to a second clipboard maintained by said second operating system,

wherein the policy data allows the data to be written to the second clipboard only if a policy defined by the policy data determines that the data was written to the first clipboard at initiation of or approved by a user, and

wherein the policy data does not allow one or more types of clipboard objects to be written to the second clipboard even a particular object of said one or more types of clipboard objects was written to the first clipboard at the initiation of or approved by said user to prevent said user from introducing a potentially hazardous type of object into said second clipboard;

the software module determining whether the data written to the first clipboard was written to the first clipboard at the initiation of or approved by the user;

upon the software module determining that the policy data does not allow the data to be written to the second clipboard because the data was not written to the first clipboard at the initiation of or approved by the user, the software module disallowing the data to be written to the second clipboard; and

upon the software module determining that the policy data allows the data to be written to the second clipboard, the software module, without human instruction, causing the data to be written to the second clipboard.

2. The one or more non-transitory machine-readable storage mediums of claim 1, wherein user input received from a physical mouse or keystrokes made on a physical keyboard will be treated as an instruction initiated by a human.

3. The one or more non-transitory machine-readable storage mediums of claim 1, wherein the instruction is deemed initiated by a human by the policy by virtue of the instruction originating from input submitted through a user interface implemented by software executing on the second operating system.

34

4. The one or more non-transitory machine-readable storage mediums of claim 1, wherein the policy data resides within the second operating system.

5. The one or more non-transitory machine-readable storage mediums of claim 1, wherein execution of the one or more sequences of instructions further causes: upon determining that a particular policy defined by the policy data considers separate data transfers to the first clipboard as a single logical unit, the software module causing each of the separate data transfers to be written to the second clipboard as one logical unit after determining the policy data allows the separate data transfers to be written to the second clipboard.

6. The one or more non-transitory machine-readable storage mediums of claim 1, wherein the software module consulting the policy data further comprises: the software module determining whether the policy defined by the policy data permits an object type of data written to said first clipboard to be written to the second clipboard maintained by said second operating system.

7. One or more non-transitory machine-readable storage mediums storing one or more sequences of instructions, which when executed by one or more processors, causes:

upon a software module, executing in a first operating environment, determining that data was written to a first bounded area of memory belonging to the first operating environment, the software module consulting policy data to determine whether, without ceasing execution and without changing priority of one or more of said first operating environment and a second operating environment, the data should additionally be written to a second bounded area of memory belonging to the second operating environment distinct from the first operating environment,

wherein said second operating environment executes concurrently with said first operating environment on a single computer,

wherein the policy data allows the data to be written to the second bounded area of memory only if a policy defined by the policy data determines that the data was written to the first bounded area of memory at initiation of or approved by a user;

the software module determining whether the data written to the first bounded area of memory was written to the first bounded area of memory at the initiation of or approved by the user; upon the software module determining that the policy data does not allow the data to be written to the second bounded area of memory because the data was not written to the first bounded area of memory at the initiation of or approved by the user, the software module disallowing the data to be written to the second bounded area of memory; and

upon the software module determining that the policy data allows the data to be written to the second bounded area of memory, the software module, without human instruction, causing the data to be written to the second bounded area of memory while the first operating environment continues to execute upon said single computer.

8. An apparatus, comprising:

one or more processors; and

one or more non-transitory computer readable storage mediums storing one or more sequences of instructions, which when executed, cause:

upon a software module, executing on a first operating system, determining that the first operating system has caused data to be written to a first clipboard

35

maintained by the first operating system, the software module consulting policy data to determine whether, without ceasing execution and without changing priority of one or more of said first operating system and a second operating system, the data should additionally be written to a second clipboard maintained by the second operating system, wherein said second operating system executes concurrently with said first operating system on a single computer, wherein the policy data allows the data to be written to the second clipboard only if a policy defined by the policy data determines that the data was written to the first clipboard at initiation of or approved by a user; the software module determining whether the data written to the first clipboard was written to the first clipboard at the initiation of or approved by the user; upon the software module determining that the policy data does not allow the data to be written to the second clipboard because the data was not written to the first clipboard at the initiation of or approved by the user, the software module disallowing the data to be written to the second clipboard; and upon the software module determining that the policy data allows the data to be written to the second clipboard, the software module, without human instruction, causing the data to be written to the second clipboard while the first operating system continues to execute upon said single computer.

9. The apparatus of claim 8, wherein user input received from a physical mouse or keystrokes made on a physical keyboard will be treated as an instruction initiated by a human.

10. The apparatus of claim 8, wherein the instruction is deemed initiated by a human by the policy by virtue of the instruction originating from input submitted through a user interface implemented by software executing on the second operating system.

11. The apparatus of claim 8, wherein the policy data resides within the second operating system.

12. The apparatus of claim 8, wherein execution of the one or more sequences of instructions further causes: upon determining that a particular policy defined by the policy data considers separate data transfers to the first clipboard as a single logical unit, the software module causing each of the separate data transfers to be written to the second clipboard as one logical unit after determining the policy data allows the separate data transfers to be written to the second clipboard.

13. The apparatus of claim 8, wherein the software module consulting the policy data further comprises: the software module determining whether the policy defined by the policy data permits an object type of data written to the first clipboard to be written to the second clipboard maintained by the second operating system.

14. An apparatus, comprising:
 one or more processors; and
 one or more non-transitory computer readable storage mediums storing one or more sequences of instructions, which when executed, cause:
 upon a software module, executing in a first operating environment, determining that data was written to a first bounded area of memory belonging to the first operating environment, the software module consulting policy data to determine whether, without ceasing execution and without changing priority of one

36

or more of said first operating environment and a second operating environment, the data should additionally be written to a second bounded area of memory belonging to the second operating environment distinct from the first operating environment, wherein said second operating environment executes concurrently with said first operating environment on a single computer, wherein the policy data allows the data to be written to the second bounded area of memory only if a policy defined by the policy data determines that the data was written to the first bounded area of memory at initiation of or approved by a user; the software module determining whether the data written to the first bounded area of memory was written to the first bounded area of memory at the initiation of or approved by the user; upon the software module determining that the policy data does not allow the data to be written to the second bounded area of memory because the data was not written to the first bounded area of memory at the initiation of or approved by the user, the software module disallowing the data to be written to the second bounded area of memory; and upon the software module determining that the policy data allows the data to be written to the second bounded area of memory, the software module, without human instruction, causing the data to be written to the second bounded area of memory while the first operating environment continues to execute upon said single computer.

15. A method for updating a software implemented in software, comprising:
 upon a software module, executing on a first operating system, determining that the first operating system has caused data to be written to a first clipboard maintained by the first operating system, the software module consulting policy data to determine whether, without ceasing execution and without changing priority of one or more of said first operating system and a second operating system, the data should additionally be written to a second clipboard maintained by the second operating system, wherein said second operating system executes concurrently with said first operating system on a single computer, wherein the policy data allows the data to be written to the second clipboard only if a policy defined by the policy data determines that the data was written to the first clipboard at initiation of or approved by a user; the software module determining whether the data written to the first clipboard was written to the first clipboard at the initiation of or approved by the user; upon the software module determining that the policy data does not allow the data to be written to the second clipboard because the data was not written to the first clipboard at the initiation of or approved by the user, the software module disallowing the data to be written to the second clipboard; and upon the software module determining that the policy data allows the data to be written to the second clipboard, the software module, without human instruction, causing the data to be written to the second clipboard while the first operating system continues to execute upon said single computer.

37

16. The method of claim 15, wherein user input received from a physical mouse or keystrokes made on a physical keyboard will be treated as an instruction initiated by a human.

17. The method of claim 15, wherein the instruction is deemed initiated by a human by the policy by virtue of the instruction originating from input submitted through a user interface implemented by software executing on the second operating system.

18. The method of claim 15, wherein the policy data resides within the second operating system.

19. The method of claim 15, further comprising: upon determining that a particular policy defined by the policy data considers separate data transfers to the first clipboard as a single logical unit, the software module causing each of the separate data transfers to be written to the second clipboard as one logical unit after determining the policy data allows the separate data transfers to be written to the second clipboard.

20. The method of claim 15, wherein the software module consulting the policy data further comprises: the software module determining whether the policy defined by the policy data permits an object type of data written to the first clipboard to be written to the second clipboard maintained by the second operating system.

21. A method for updating a bounded area of memory of a computer system, comprising:

upon a software module, executing in a first operating environment, determining that data was to be written to a first bounded area of memory belonging to the first operating environment, the software module consulting policy data to determine whether, without ceasing execution and without changing priority of one or more

38

of said first operating environment and a second operating environment, the data should additionally be written to a second bounded area of memory belonging to the second operating environment distinct from the first operating environment,

wherein said second operating environment executes concurrently with said first operating environment on a single computer,

wherein the policy data allows the data to be written to the second bounded area of memory only if a policy defined by the policy data determines that the data was written to the first bounded area of memory at initiation of or approved by a user;

the software module determining whether the data written to the first bounded area of memory was written to the first bounded area of memory at the initiation of or approved by the user;

upon the software module determining that the policy data does not allow the data to be written to the second bounded area of memory because the data was not written to the first bounded area of memory at the initiation of or approved by the user, the software module disallowing the data to be written to the second bounded area of memory; and

upon the software module determining that the policy data allows the data to be written to the second bounded area of memory, the software module, without human instruction, causing the data to be written to the second bounded area of memory while the first operating environment continues to execute upon said single computer.

* * * * *